RL-TR-95-227
Final Technical Report
November 1995

DATA FUSION: PERFORMANCE ENHANCEMENT VIA RESOLUTION CELL PROCESSING

University of Connecticut

Constantino Rago and Peter Willett

19960304 081

**Rome Laboratory
Air Force Materiel Command
Rome, New York**

DTIC QUALITY INSPECTED 1

This report has been reviewed by the Rome Laboratory Public Affairs Office (PA) and is releasable to the National Technical Information Service (NTIS). At NTIS, it will be releasable to the general public, including foreign nations.

RL-TR-95- 227 has been reviewed and is approved for publication.

APPROVED: *Mark G. Alford*

MARK G. ALFORD
Project Engineer

FOR THE COMMANDER: *Gary D. Barmore*

GARY D. BARMORE, Major, USAF
Deputy Director of Surveillance & Photonics

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. AGENCY USE ONLY (Leave Blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	November 1995	Final Jul 93 – Aug 95

4. TITLE AND SUBTITLE

DATA FUSION: PERFORMANCE ENHANCEMENT VIA RESOLUTION CELL PROCESSING

5. FUNDING NUMBERS

C – F30602-94-C-0060
PE – 61102F
PR – 2304
TA – E8
WU – PA

6. AUTHOR(S)

Constantino Rago and Peter Willett

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

University of Connecticut
U-157
Storrs CT 06269

8. PERFORMING ORGANIZATION REPORT NUMBER

N/A

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Rome Laboratory/OCSM
26 Electronic Pky
Rome NY 13441-4514

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

RL-TR-95-227

11. SUPPLEMENTARY NOTES

Rome Laboratory Project Engineer: Mark G. Alford/OCSM/(315) 330-3573

12a. DISTRIBUTION/AVAILABILITY STATEMENT

Approved for public release; distribution unlimited.

12b. DISTRIBUTION CODE

13. ABSTRACT (Maximum 200 words)

This report discusses the resolution cell grid topological connectedness as understood by a plurality of sensors. This is a data fusion obstacle ignored by previous multisensor fusion treatments. This connectedness, if properly considered, can provide enhanced resolution. The noncoherent case draws attention to the sort of processing necessary when multiplatform fusion is used. The problem is described and analyzed for both the static and dynamic cases. Since performance gains depend on the fusion rule used, pre-detection fusion rules are developed. "Same-sensor" fusion makes use of the complementary features, detection, and resolution of long and short pulses. This teaming approach appears very promising. If different single sensor pulses are fused, then coherent processing is possible. A pair of different complementary pulses is transmitted abutted in the time domain and the returns optimally fused. Combined constant-frequency and linearly frequency-modulated pulses are treated, and order magnitude tracking error reduction over isolated single pulse performance is demonstrated.

14. SUBJECT TERMS

Data fusion, Sensor parameter enhancement, Resolution cell processing

15. NUMBER OF PAGES

76

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	UL

Contents

List of Figures

Chapter 1

The Noncoherent Case

Abstract

If members of a suite of sensors from which fusion is to be carried out are not co-located, it is unreasonable to assume that they share a common resolution cell grid; this is generally ignored in the data fusion community. In this report we explore the effects of such "noncoincidence", and we find that what at first seems to be a problem can in fact be exploited. The idea is that a target is known to be confined to an intersection of overlapping resolution cells, and this overlap is generally small.

We examine noncoincidence from two viewpoints: tracking and detection. With respect to tracking our analysis is first *static*, by which is meant that we establish the decrease in measurement error; and then *dynamic*, meaning that the overall effect in the tracking problem is quantified. The detection viewpoint considers noncoincidence as it has impact on a pre-detection fusion system. Specifically, the role of the fusion rule is examined, and the use of noncoincidence to improve *detection* performance (rather than that of tracking) is explored. Finally, we look into the idea of "single-sensor" noncoincidence.

1

1.1 Introduction

Observations from a suite of two or more radars can be fused at three levels: as tracks, as detections, and as observations. In the first, each sensor performs its own tracking, with the resulting estimated target trajectories and covariances combined. In the second, all sensor detections are combined as input to a single tracking algorithm. It is common to refer to these as *track fusion* and *post-detection fusion*, and schemes for each are available in [9] and its references. In the third approach, often called *pre-detection fusion*, target reports from the individual sensors are combined and evaluated prior to being passed to a tracking routine – an excellent treatment and bibliography of the topic is in [2]. The distinction between the second and third approaches is in the location of the hit/no-hit decision-making: in the former this is local, and in the latter it is fused.

The three schemes are hierarchical in the location of fusion within the data-processing chain, and it is apparent that the pre-detection case, with its fusion of the "rawest" of the data, should offer the best performance. There are challenges, however:

- The resolution cell "grids" from the various sensors in general do not line up with each other, and hence it is not immediately clear how reports from resolution cells that overlap but do not coincide should be fused.

- Unless particular care is taken, the sensors will not be time-synchronized. As such, due to target motion, reports from the same target can appear in resolution cells which do not overlap.

These problems are daunting, and perhaps explain the concentration of algorithmic results on the upper parts of the hierarchy, and the parallel abundance of theoretical results on decentralized decision-making among *common* hypotheses (*i.e.* all sensors test the same resolution cell at the same time).

It is our intention, in this report, to examine the *first* of the above problems, that of resolution-cell disagreement. We shall show, in fact, that the non-coincidence among grids amounts, somewhat surprisingly, to *enhanced* resolution: making proper use of overlap can actually improve performance. To do so we must ignore the second problem: synchronicity

among sensors' observations will be assumed. (Relaxation of this will be the subject of future research.)

We are concerned that the discussion may alienate practically-minded readers at this point: it is not realistic that each member of a suite of sensors has its attention focused on the same target at the same time. With reference to Figure 1.1, if the two near sensors are assumed to scan in lock-step and are surveying a comparatively distant volume, it is at least feasible. But sensors are seldom located so close together, and in any case cannot be relied upon for such accurate synchronization, so to dwell on such a case would be an apology. Our point is that there are major obstacles to the practical deployment of a pre-detection fusion system, obstacles that have not, to our knowledge, previously received analytical attention; and our goal is to examine *one* of them.

At any rate, let us return to Figure 1.1. It is clear that the resolution cells from the sensors do not overlap precisely, and hence it is not surprising to find in Figure 1.2 that a single target results in a confusing collection of detection information at the fusion center.

The disagreement in resolution-cell overlap appears at first to be a drawback. However, a closer examination of the figure reveals that noncoincidence may actually be a *feature*, in that the target is known to be confined to the *overlap* region between the two cells with detections. There are two obvious choices:

Measurement-based fusion (MBF). The centroid over all resolution cells containing detections ("hits") corresponding to the target being tracked is the measurement conveyed to the tracker. With similar-sized resolution-cells this amounts to the centroid of the centroids (of the individual resolution cells); the measurement variance is scaled by the reciprocal of the number of sensors recording hits. (This is specific to the case that the sensors' measurement covariances are identical; more generally, a scaling according to the familiar "parallel resistors" law is appropriate.) This fusion scheme does not exploit the resolution cell overlap geometry, and hence we shall use this as a benchmark for comparison.

Resolution-cell-based fusion (RCBF). The measurements conveyed to the tracker are the centroids of the resolution cell overlaps. The measurement error covariances

3

correspond to the sizes of the cell-overlaps. This scheme *does* exploit the overlap geometry.

With reference to the two-sensor example shown in Figure 1.2, note that each true position (denoted by a o) produces a hit in two different resolution cells, one for each sensor. Each of these resolution cells has a centroid (denoted by an ×), and the average of these (denoted by a +) is taken as the measurement by MBF. In each case, however, the true position lies within the *overlap* region of the two resolution cells, and the centroid of this is what is used as a measurement in RCBF. In this two-sensor case, the MBF measurement and the RCBF measurement are indistinguishable (although this is not always the case). However, note that when the overlap is small (for example, in the fourth snapshot of Figure 1.2), RCBF has the opportunity to "tell" the tracker that the measurement being supplied is very accurate.

In the "Tracking" section we shall perform a static probabilistic analysis of the overlap, and from this shall infer the statistics of the measurement noise under the two schemes. By "static" we mean that no tracking is involved – our goal here is simply to evaluate the relative quality of the measurements. It turns out that the main difference between MBF and RCBF is not in the measurements, which tend to be very close; but rather in the measurement noise variance to be taken as input to the Kalman filter. RCBF supplies the fact that this variance is time-varying, while MBF does not. We then turn to the dynamic problem, and analyze the effect of a time-varying measurement variance, and shall observe that a stochastic measurement variance (with mean $\bar{R}$) is more amenable to estimation than one in which the measurement variance is deterministic and fixed at $\bar{R}$. We conclude the section with a Monte-Carlo comparison of MBF to RCBF, and shall observe that RCBF is indeed superior, although by a perhaps not-overwhelming amount.

In the "Detection" section we explore in the context of noncoincidence the use of various fusion rules for detection, and show that the rule chosen must be a compromise between detection and estimation performance. We then look at the same problem from a different point of view, and ask: by how much can detection in RCBF be improved over that in MBF *when both have the same tracking performance*? The "trick" here will be to use a longer pulse in RCBF such that tracking performance is identical to MBF, but SNR is increased.

Finally, we explore *deliberate* noncoincidence, in which a single waveform-agile sensor has measurements from long and short pulses combined optimally in the pre-detection sense. In this section there is no problem with an assumption of synchronization.

1.2 Tracking

1.2.1 Analysis of the static problem

With reference again to Figure 1.2, it is clear that the size of the overlap, and hence the measurement error covariance matrix, is not constant. In this section we treat it as random, and our goal, naturally, is to develop an expression for its probability density function. It must be noted that, although in the situation of Figure 1.2 there is hope of doing this exactly, in more general situations of non-congruent resolution cells (or even congruent ones with different orientations), or with more than two sensors, an exact solution is unlikely Our approach here will be approximate: we shall work in only one dimension (a random length), with the idea that results in higher dimensions (random areas or volumes) are obtainable via products of these random lengths. It will turn out that the one-dimensional analysis is straightforward and crisp; equally important, it will turn out that the extension-via-product idea is reasonably accurate.

At any rate, let us begin by examining the two-sensors (and, of course, unidimensional) case. Given the position of a target there are two intervals (one for each sensor) that are known to contain it. Our approach is to fix one of the intervals, and allow the target's position and the other interval to vary – see the Appendix for the derivation. We find that the probability density function of the overlap interval (Δ) is given by:

$$f_2(\Delta) = \frac{2}{L_1 L_2}\Delta + \frac{L_2 - L_1}{L_2}\delta(\Delta - L_1) \tag{1.1}$$

for $0 \leq \Delta \leq L_1$, where L_1 and L_2 are the lengths of the resolution cells ($L_2 > L_1$ without loss of generality). Conditioned on Δ the target's location is uniformly distributed along the overlap interval, and hence the estimation error has variance $\sigma^2 = \Delta^2/12$. When $L_1 = L_2 \equiv L$ (the situation of interest for two similar radars and a distant target), it

5

can be shown, via removing the conditioning on Δ, that σ^2 is uniformly distributed on $[0, L^2/12]$; that is,

$$f(\frac{\Delta^2}{12}) = f_{\sigma^2}(\sigma^2) \;\; = \;\; 12/\sigma^2 \tag{1.2}$$

for $0 \le \sigma^2 \le L^2/12$ and zero elsewhere. For three sensors the probability density function of the overlap interval is given by:

$$f_3(\Delta) \;\; = \;\; 2\Delta \left(\frac{L_1 + L_2 + L_3}{L_1 L_2 L_3} - \frac{3\Delta}{L_1 L_2 L_3} \right) + \frac{(L_3 - L_1)(L_2 - L_1)}{L_2 L_3} \delta(\Delta - L_1) \tag{1.3}$$

for $0 \le \Delta \le L_1$, where $L_3 > L_2 > L_1$ are the lengths of the resolution cells. The variance of the estimation error (again when $L_1 = L_2 = L_3 \equiv L$) is given by

$$f(\frac{\Delta^2}{12}) = f_{\sigma^2}(\sigma^2) \;\; = \;\; \frac{36}{L^2} \left(1 - \frac{\sqrt{12\sigma^2}}{L} \right) \tag{1.4}$$

for $0 \le \sigma^2 \le L^2/12$ and zero elsewhere. Extension to more than three sensors is straightforward, but we do not report it here.

Let us now examine the worth of our analysis in the multidimensional case. This we must do by example, and the situation we shall deal with is of two radars separated by 3.2 km, each with a resolution in range of 100 meters and a resolution in azimuth of 1 degree. We have numerically evaluated measurement variances both in range (R) and in azimuth (θ) for a set of R values, by which we mean that we have used a computer to establish the appropriate resolution cell grids and have calculated the average variances encountered while traversing circles (centered between the sensors) at various ranges R. We compare the results to our analysis, from which we have in range and azimuth

$$\sigma_r^2 = \delta_r^2/24 \qquad \sigma_\theta^2 = (\delta_\theta R)^2/24 \tag{1.5}$$

respectively, where δ_r and δ_θ are the range and azimuth resolution, and R is the range.

The results are given in Figures 1.3 and 1.4. We observe a good match between the predicted and computed values, with the exception of the variance in the range axis for small values of R; the reason is that our assumption of on-average square resolution cells is not valid here.

6

Our goal is to compare RCBF and MBF: how different are they? With regard to the measurements themselves, the answer is: not very. It is easy to show that the measurements obtained in the unidimensional case by RCBF and MBF are identical. In more than one dimension this is not always precisely true, but it is at least approximately so when the geometry (*i.e.* shape and rotation) of the various resolution cells is similar, as it is in our case of a distant target and proximate sensors. What we have dealt with in this section is the measurement *covariance*, for it is here that we expect there to be differences – to be specific, RCBF supplies to the tracker the correct covariance, while MBF does not.

Consider a one-dimensional situation: with a single sensor (and an assumed uniformity), the measurement error is $L^2/12$, with L the length of the resolution cell. Using MBF for the measurement, there is little choice but to assume an error variance of $L^2/24$ for two sensors and $L^2/36$ for three. If we instead use RCBF for the measurement, we get error variances of $L^2/24$ and $L^2/40$ for two and three sensors, respectively. Now, these are identical for two sensors and quite close for three: it does not appear, in the average sense, that RCBF offers a noticeable advantages over MBF. However, averages do not reveal the true situation: *the measurement error variance itself is stochastic and varying*, and this variation can be significant. We shall explore its effect in the next section

1.2.2   Analysis of the dynamic problem

As we have seen, the measurement noise covariance for RCBF is a stochastic process. This implies that the evolution of the covariance matrix of the estimated states (using a Kalman filter) is not given by the solution of the deterministic Riccati equation, but rather by the solution of its stochastic generalization.

The state covariance matrix is updated by the Kalman filter as follows:

$$\mathbf{P}_{k+1|k} = \mathbf{F}\mathbf{P}_{k|k}\mathbf{F}' + \mathbf{Q}$$

$$\mathbf{S}_{k+1} = \mathbf{H}\mathbf{P}_{k+1|k}\mathbf{H}' + \mathbf{R}_k$$

$$\mathbf{W}_{k+1} = \mathbf{P}_{k+1|k}\mathbf{H}'\mathbf{S}_{k+1}^{-1}$$

$$\mathbf{P}_{k+1|k+1} = \mathbf{P}_{k+1|k} - \mathbf{W}_{k+1}\mathbf{S}_{k+1}\mathbf{W}'_{k+1} \tag{1.6}$$

In the deterministic case, the covariance evolves to a steady state which characterizes the performance of the filter, and this solution can be computed from the associated algebraic Riccati equation (ARE). When the measurement noise has a stochastic variance ($\mathbf{R}_k$), as in our case, there *is* no steady-state. However, it is legitimate to seek the *expected value of the error-covariance* in this case, and that is the subject of this section.

Let us assume that the measurement noise covariance matrix can be modeled as:

$$\mathbf{R}_k = \mathbf{R}_o r(k) \tag{1.7}$$

where $r(k)$ is an *iid* one-dimensional stochastic process that scales the covariance matrix $\mathbf{R}_o$ to take into account the overlap.

The method here proposed is based on the HYCA (hybrid conditional average) method of [8], and involves discretizing $r(k)$ into M levels ($r(k) \in \{r_i\}_{i=1}^M$), where each level has probability p_i. Let us denote the covariance matrix associated with level i as $\mathbf{P}_i(k|k)$. We update each of these $\mathbf{P}_i(k|k)$ with the measurement noise covariance associated with level r_j, resulting in M covariance matrices $\mathbf{P}_{ij}(k+1|k+1)$; the updated covariance matrix for level j is then given by

$$\mathbf{P}_j(k+1|k+1) = \sum_{i=1}^M \mathbf{P}_{ij}(k|k)p_i \tag{1.8}$$

The matrices $\mathbf{P}_{ij}(k|k)$, which reflect the state error covariance given that $r(k-1) = r_i$ and $r(k) = r_j$, are obtainable from (1.6). The mean state covariance matrix (*i.e.* what we want) is given by

$$\mathbf{P}(k|k) = \sum_{j=1}^M \mathbf{P}_j(k|k)p_j \tag{1.9}$$

evaluable when this matrix Markov process reaches steady state. Figure 1.5 shows a flow-diagram that reflects the mechanics of the iterations.

In Figure 1.6 we have compared the actual mean state covariance matrix elements from 50 Monte Carlo runs both with the mean obtained using our method and with that obtained via an ARE (*i.e.* that obtained under the assumption that measurement noise variance is constant and equal to its mean). For this case we have used the two-dimensional kinematic model of equation (1.10) – the same as used to generate Figure 1.2, and the same to be studied more extensively in the following section – with three sensors and resolution

8

cells of side $L = 1000m$. Perhaps the most obvious conclusion is that the ARE-based solution is pessimistic. Heuristically, when the measurement noise is time-varying the occasional incidents of very small measurement noise (which is *known* to be very small) allow a Kalman filter to "zero in" very accurately on its target's state. Again heuristically, a constant measurement noise variance is in a sense a worst case. A second conclusion is that our analysis appears to be quite good.

1.2.3 Results for the dynamic problem

Our application is not the static but rather the dynamic estimation problem, *i.e.* tracking. Here the analysis becomes complicated, and it is very difficult to make a fair rigorous comparison between a Kalman filter using the average of centroids as measurements and a *fixed* (and statistically incorrect) measurement covariance (MBF), with one using the centroid of the overlap as the measurement and an appropriate *time-varying* measurement covariance (RCBF). Our goal is, of course, to make some statements about the relative performance of MBF and RCBF. To do so we need ground-rule assumptions, and these will be:

- A single target is being tracked by a collection of like sensors, from each of which range and azimuth information is available for all detections. (There is no data association problem here, nor is there the complication of deciding how to fuse doppler information.)

- The resolution cell grids of the various sensors do not overlap precisely, but are known (and hence the overlap region is available to the fusion center).

- For all sensors the false-alarm probability is zero and the probability of detection is unity. (This assumption is of course unrealistic, but it is clear that the inclusion of practical detection models will give results more favorable to the resolution-cell-based paradigm. The detection aspects of non-coincidence are analyzed in the next section.)

Here we present a comparison based on a Monte Carlo simulation of a two-dimensional, two-sensor scheme tracking a target that obeys a second order kinematic model (actually

9

the same model used to create the trajectory of Figure 1.2) given by

$$x_{k+1} = Ax_k + Bv_k \qquad (1.10)$$

where

$$A = \begin{bmatrix} 1 & T & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & T \\ 0 & 0 & 0 & 1 \end{bmatrix} \qquad B = \begin{bmatrix} T^2/2 & T & 0 & 0 \\ 0 & 0 & T^2/2 & T \end{bmatrix}$$

and where the process noise variance $E\{v_k^2\} = .1$, and the sampling period (T) is 20 sec. Note that there is no "measurement" equation; the measurements are given either by the centroid of the cell-overlap (RCBF) or the centroid-of-centroids (MBF) – recall that our simplifying assumption was that the probability of detection is unity and the probability of false-alarm is zero. The measurement noise covariance matrix is proportional to the overlap area in the first case, and is constant in the second case.

We now compare the performances of these two techniques via Monte Carlo simulation. Those which we present are based on 50 Monte-Carlo runs; each trajectory has 50 points in time, and the errors are averaged along each trajectory and run.

Figure 1.7 and Figure 1.8 each show errors for RCBF and for MBF, these errors being of position and of velocity, respectively. In these plots the abscissae represent inter-sensor distance – what we mean to show here is that the situation changes as sensors move apart from one another, due to the variation in "disorder" of the resolution-cell grids. First, note that RCBF offers an improvement in tracking accuracy over MBF – this improvement is not constant, but is generally of the order of 10%. Second, note that the improvement becomes more marked as the inter-sensor separation increases – again, this is due to the increasing geometric "disorder" obtained when the resolution-cell grids emanate from further-separated points. At any rate, we note that RCBF offers improvement over MBF even in the two-sensor case, *despite there having been no such improvement for the static estimation.* However, we must admit that the gains are modest.

We further explore the effect of the *number* of sensors. To be as fair as possible, the inter-sensor distance is kept fixed: for two they are d meters apart; for three in a triangle

10

of side d; for four in a square of side d; and for five in a square of side $\sqrt{2}d$ with the fifth in the center. From Figure 1.9 it is clear that the improvement of RCBF over MBF becomes more marked as the number of sensors increases.

1.3 Detection

The previous analysis was entirely concerned with the effects of noncoincidence on estimation and on tracking. Detection was ignored, insofar as everything was based on perfect (zero probability of false-alarm and of miss) detection.

Perhaps most central to a pre-detection fusion system is the fusion *rule*. This describes explicitly what combinations of sensor outputs (binary is assumed here) result in a detection. An example of such is the *or* rule, under which a detection is declared if at least one of the sensors reports a hit. The *or* rule is a legitimate fusion strategy; yet care must be taken when it is applied in the context of noncoincidence, since the concept of "overlap" when only one sensor reports a hit is not well-defined.

To proceed further we need not only to understand fusion rules, but also to understand their effect on measurement errors in RCBF fusion. These were investigated in detail in [6] and we report from those findings next. We shall use these results to predict which fusion rules are best to be used in tandem with RCBF fusion, and shall further explore a scheme whereby RCBF fusion is used to improve detection performance.

1.3.1 Fusion rules

For noncoincidence to be useful, some kind of time-synchronization between the sensors is necessary, and as described in [6] this is tantamount to similar and nearby sensors. What this means for us here, assuming binary decision rules at the sensor level, is that at the fusion center the optimum rules must be of the k-out-of-N variety: declare a detection if and only if at least k of the N sensors report hits. There are two measures of performance which interest us: the global probability of detection (this is the obvious one) and the average measurement error variance. Intuition suggests that the larger k, the better the

performance with respect to the latter; however, unfortunately, performance with respect to the former is often best for small k. Our goal in this section is to make some statements about this trade-off.

Using the same threshold at each sensor, the probability of false alarm (P_F) and the probability of detection (P_D) at the fusion center are given by [3]

$$P_F = \sum_{n=k}^{N} \left(\begin{array}{c} N \\ n \end{array} \right) P_{fi}^n (1 - P_{fi})^{N-n} \tag{1.11}$$

$$P_D = \sum_{n=k}^{N} \left(\begin{array}{c} N \\ n \end{array} \right) P_{di}^n (1 - P_{di})^{N-n} \tag{1.12}$$

where k defines the fusion rule as above, and where P_{fi} and P_{di} are the probabilities of false alarm and of detection at the sensor level. These two quantities depend on the local threshold and statistics under the hypotheses, but it should be noted that with k fixed this "operating point" is an explicit function of P_F — the only optimization needed is that over k.

Again assuming a k-out-of-N rule, we have the global measurement error variance in each dimension (azimuth and range)

$$\epsilon^2 = \sum_{n=k}^{N} \left(\begin{array}{c} N \\ n \end{array} \right) (\epsilon_n^2) P_{di}^n (1 - P_{di})^{N-n} \tag{1.13}$$

where ϵ_n^2 is the measurement error variance given that n sensors report hits. This quantity is reported in the previous section, and as the first row in Table 1.1.

	$n = 2$	$n = 3$	$n = 4$	$n = 5$
ϵ_{ns}^2	$L^2/24$	$L^2/40$	$L^2/60$	$L^2/84$
ϵ_{nd}^2		$.9\epsilon_{ns}$		

Table 1.1: Expected value of measurement errors per dimension when n sensors report hits, and where the length of the resolution cell is L. The first row reports "static" (s) results from section 2.1; the second row gives "dynamic" (d) results.

There is a compromise between the total probability of detection and the mean error measurement variance. The largest reduction in the latter happens for the *and* rule; on

the other hand, the highest probability of detection may correspond to some intermediate rule. To illustrate this we have computed the above quantities for the CA-CFAR scenario in which the sensors' observations are governed by [7]

$$
\begin{aligned}
p(X_i > x|H) &= (1+x)^{-m} \\
p(X_i > x|K) &= \left(1 + \frac{x}{1+S}\right)^{-m}
\end{aligned}
\tag{1.14}
$$

under the target-absent (H) and target-present (K) hypotheses, respectively, where S represents the signal to noise ratio. In Figure 1.10 we see the total probability of detection and the improvement in the mean measurement variance error $(\epsilon^2_{RCBF} - \epsilon^2_{MBF}) * 100/\epsilon^2_{MBF}$. It is well-known that for the CA-CFAR scenario and reasonable SNR and P_F values the *and* rule is poor in terms of detection. From the figure, however, we observe that the performance in terms of measurement error is poor for low k – as such, it appears that an intermediate value of k is a good compromise, perhaps $k = 3$ or 4 for the $N = 5$ case given, with similar results for other values of N.

1.3.2   Better detection through overlap

It is clear that we can obtain certain measurement noise reduction by exploiting the overlap between the resolution cells. But there is an alternative: we can enlarge the sensors' resolution cells in the RCBF fusion scheme, and depend upon the overlap feature to reduce it again. The purpose is the following: a larger resolution cell, at least in the range direction, is the result of a longer pulse, and a longer pulse means a higher SNR.

The natural question is by what factor the pulse length can be increased such that the RCBF and MBF fusion schemes enjoy the same measurement errors. Section 1.2.1 and the first row of Table 1.1 provide one answer. The results of section 1.2.3 indicate, however, that this is not the complete answer: it was found that the ability of RCBF fusion to use the stochastic (varying) measurement covariance (*i.e.* to provide the tracker with an estimate of its current measurement's accuracy, which will be good when the overlap region is small) is its true strength. A conservative (actually a very conservative) figure for the reduction in estimation error is 10%. It must be acknowledged that: this figure is an approximation; that

there is not a linear correspondence between measurement and estimation accuracies; and that the analysis of section 1.2.3 was based on ideal detection performance. Nonetheless, our analysis appears to be at least approximately correct, and in Figure 1.11 we can see the comparison between RCBF (with both static- and dynamic-based pulse-length expansion factors) and MBF fusion for a fixed probability of false alarm. Similar results, but for a fixed probability of detection are shown in Figure 1.12. The decrease in miss probability in the first case or probability of false alarm can be significant.

1.4   The Single-Sensor Case

Until now we have worked with non-colocated radars. Here we apply some of the previous ideas to the single radar case — this is in some sense more appealing than the multi-sensor scenario, since there is no problem of time synchronization. At any rate, we concentrate on what happens when we introduce a special kind of non-coincidence to a single radar.

Let us assume that the radar is *waveform-agile* [7] to the extent that it has the capability of using a long pulse as well as a short pulse (see Figure 1.13). Such agility is commonly applied to deal with targets which are either very near to the sensor (where a short pulse is required for resolution) or distant (where a long pulse is necessary for detection). Rather than use them for such separate ends, here we shall examine the case that *all space is illuminated by both long and short pulses, and the respective information from these is fused prior to detection.*

Associated with a long pulse is good detection performance but poor range resolution; conversely a short pulse gives a good range resolution but a poor probability of detection. Based on this, we analyze the following scheme to combine these two measurements: use the long pulse return to detect, and then estimate range via the short pulses. The detection problem is a standard one, and therefore we will use a likelihood ratio with a CA-CFAR model for the hypotheses. The range estimation problem using the short pulse returns is more interesting. In the next section we will derive the optimum (in the minimum mean square error sense) range estimator for this situation.

14

1.4.1 The optimum range estimator

Let us assume that the duration of the long pulse is N times the duration of the short one. Then, for each "long" cell marked with a return we have N "short" cells; an associated return vector is, for example

$$\mathbf{u} = \underbrace{[0010001 \cdots 001]}_{N} \tag{1.15}$$

where the 1's represent a return from a short cell which is above a threshold.

The probability of any of these return vectors $\mathbf{u}$ conditioned on the fact that the actual target is in cell j (*i.e.* the actual range is R, which is in cell j) is given by

$$p(\mathbf{u}|R) = \begin{cases} P_1 = P_d P_f^{s-1}(1-P_f)^{N-s} & ; u_j = 1 \\[2mm] P_0 = (1-P_d)P_f^s(1-P_f)^{N-s-1} & ; u_j = 0 \end{cases}$$

where p denotes a probability mass function, P_d, P_f are the probabilities of detection and false alarm associated with the short pulse and $s = \sum_{i=1}^{N} u_i$.

The minimum mean square error estimate (MMSE) of the range, $\hat{R}(\mathbf{u})$, is given by

$$\hat{R}(\mathbf{u}) = E(R|\mathbf{u}) = \int Rp(R|\mathbf{u}) = \int R \frac{p(\mathbf{u}|R)p(R)}{p(\mathbf{u})} dR \tag{1.16}$$

where we will assume R is uniformly distributed over the length (L_1) of the range resolution cell corresponding to the long pulse. With this we have that (1.16) can be rewritten as

$$\hat{R}(\mathbf{u}) = \frac{1}{L_1} \sum_{i=1}^{N-1} \int_{iR/N}^{(i+1)R/N} R \frac{p(\mathbf{u}|R)}{p(\mathbf{u})} dR \tag{1.17}$$

We need to compute $p(\mathbf{u})$,

$$p(\mathbf{u}) = \int_R p(\mathbf{u}|R)p(R)dR$$

$$= \frac{1}{L_1} \sum_{i=1}^{N-1} \int_{iR/N}^{(i+1)R/N} p(\mathbf{u}|R)dR$$

$$= \frac{1}{L_1} \left[P_1 \sum_{i=0}^{N-1} \frac{L_1}{N} u_i + P_0 \sum_{i=0}^{N-1} \frac{L_1}{N}(1-u_i) \right]$$

$$= \frac{1}{N} [sP_1 + (N-s)P_0] \tag{1.18}$$

Now we can proceed to compute the minimum mean square error estimator

$$\hat{R}(\mathbf{u}) = \frac{1}{L_1 p(\mathbf{u})} \sum_{i=0}^{N-1} \int_{iR/N}^{(i+1)R/N} R p(\mathbf{u}|R) dR$$

$$= \frac{L_1/N^2}{p(\mathbf{u})} \left[P_1 \sum_{i=0}^{N-1} \frac{2i+1}{2} u_i + P_0 \sum_{i=0}^{N-1} \frac{2i+1}{2} (1-u_i) \right] \qquad (1.19)$$

$$= \frac{L_1/N^2}{(sP_1 + (N-s)P_0)} \left[(P_1 - P_0)s(\frac{1}{2} + \overline{\mathbf{u}}) + P_0 \frac{N^2}{2} \right]$$

where $\overline{\mathbf{u}} = (1/s)\sum_{i=0}^{N-1} i u_i$

As easily seen, the MMSE estimator is linear in $\mathbf{u}$, with parameters depending on the Receiver Operating Characteristic (ROC) operating point. Its variance is given by

$$E((\hat{R}(\mathbf{u}) - R)^2) = \sum_{\text{all } u} \frac{1}{L} \sum_{i=0}^{N-1} (\hat{R}(\mathbf{u}) - iL/N)^2 \qquad (1.20)$$

1.4.2   Optimal pulse-length ratio

Here we explore the choice of N, the number of short pulses per long pulse. Based on the previous analysis, and for a given operating point, we can choose N such that the variance of the range estimator be minimum. In Figure 1.14 we have plotted the variance of the range estimator as a function of the probability of false alarm for different $N's$. We see that the optimum N depends on the desired probability of false alarm, with an optimum $N = 4$ or 5 for intermediate probabilities of false alarm, and $N = 2$ or 3 for low probabilities of false alarm. When the probability of false alarm is extreme, then the relationship between the short and long pulse becomes irrelevant, as the variance is dominated by either false alarms or missed detections, and not estimation errors.

1.5 Summary

Theoretical investigations of pre-detection fusion systems generally work from the idealized assumptions that all sensors gather data from the same phenomenon at the same time. In this report we deal with the first of these: we explore the practically-significant case that the local data relates to somewhat different phenomena, in that the constituent resolution cell grids do not coincide.

The intent of the research was to quantify the effect of what was hoped to be a very attractive feature. To this end:

- We compared two schemes. Resolution-cell-based fusion (RCBF) exploits the fact that detections from various sensors are dependent through the actual location of the underlying target. Conversely, measurement based fusion (MBF) treats detections from the various sensors as if they were independent, and simply averages them appropriately. We find, via a "static" analysis (section 2.1), that both the measurements themselves and their *average* accuracies are quite similar under the two schemes. The difference is that RCBF is able to supply (to the tracker) an accurate depiction of each measurement's *time-varying* covariance.

- To study ("dynamically") the effect of such *random* measurement covariance, we have used both analysis (section 2.2) and simulation (section 2.3). Our results indicate that RCBF *is* an improvement over MBF.

- We have investigated the effect of the detection fusion rule (section 3.1). It turns out that there is a trade-off between detection and resolution.

- We observed that the gains in resolution can be exchanged, via pulse-lengthening, for higher SNR and significantly improved detection performance (section 3.2).

- We introduced and analyzed a method whereby *deliberate* noncoincidence, using long and short pulses in tandem at a single sensor, was able to improve upon the detection performance of the former while retaining the resolution of the latter (section 4).

17

The last point is particularly promising, and may indicate a good direction for future research. As for the others, it turned out that our original optimism was unsatisfied, in that the improvements we have seen are modest. Our conclusion, however, is positive: there is certainly nothing to fear from pre-detection fusion in a realistic framework, and indeed, it works *better* than under idealized conditions.

Appendix

We will first examine the two-sensor/one-dimension case. In this case, the resolution cells are just intervals. Given the position of a target there are two intervals (one corresponding to each sensor) that contain it. We will denote the length of these intervals as L_1 and L_2, and specify $L_1 < L_2$ without loss of generality. Let us fix the position of one of the intervals (L_1); then the target's position (we will denote it by x) could be any point along the first interval. The starting point of the second interval (we will denote this random variable as y) could be any point in $[0, L_2]$. A reasonable model for x and y is assume that x is uniformly distributed in $[0, L_1]$ and y is uniformly distributed in $[0, L_2]$. The length of the overlapping interval (Δ) conditioned on a particular x is given (as a function of y) by:

- for $0 < y < x \rightarrow \Delta = L_1 - y$

- for $x < y < L_1 \rightarrow \Delta = y$

- for $L_1 < y < L_2 \rightarrow \Delta = L_1$

This relationship is shown in Figure 1.15 for the case that $x < L_1/2$ and $x > L_1/2$.

The probability density function of Δ conditioned on x can be found from the previous relationship between Δ and y and the fact that y is uniformly distributed in $[0, L_2]$

$$f(\Delta|x) = (1/L_2)(U(x) - U(L_1 - x)) + (2/L_2)(U(L_1 - x) - U(L_1)) + \frac{L_2 - L_1}{L_2}\delta(x - L_1) \quad (1.21)$$

We can now find the unconditional probability density function of Δ as

$$f(\Delta) = \int_0^{L_1} f(\Delta|x)f(x)dx \tag{1.22}$$

The solution of this integral is (1.1).

The generalization for more than two sensors is immediate. Let us consider the three-sensor case: the intersection of three intervals (which we will denote by γ) is equal to the intersection of the overlap between two of them (Δ) and the remainder one. Then

$$f(\gamma|\Delta) = \frac{2}{\Delta L_3}\gamma + \frac{L_3 - \Delta}{L_3}\delta(\gamma - \Delta) \qquad 0 \le \gamma \le \Delta \le L_3 \tag{1.23}$$

19

Now we can find the unconditional *pdf* of γ in the same way we did Δ, the only difference being that Δ has a *pdf* given by (1.1) instead of a uniform. Extension to more than three sensors is straightforward, but is not given here.
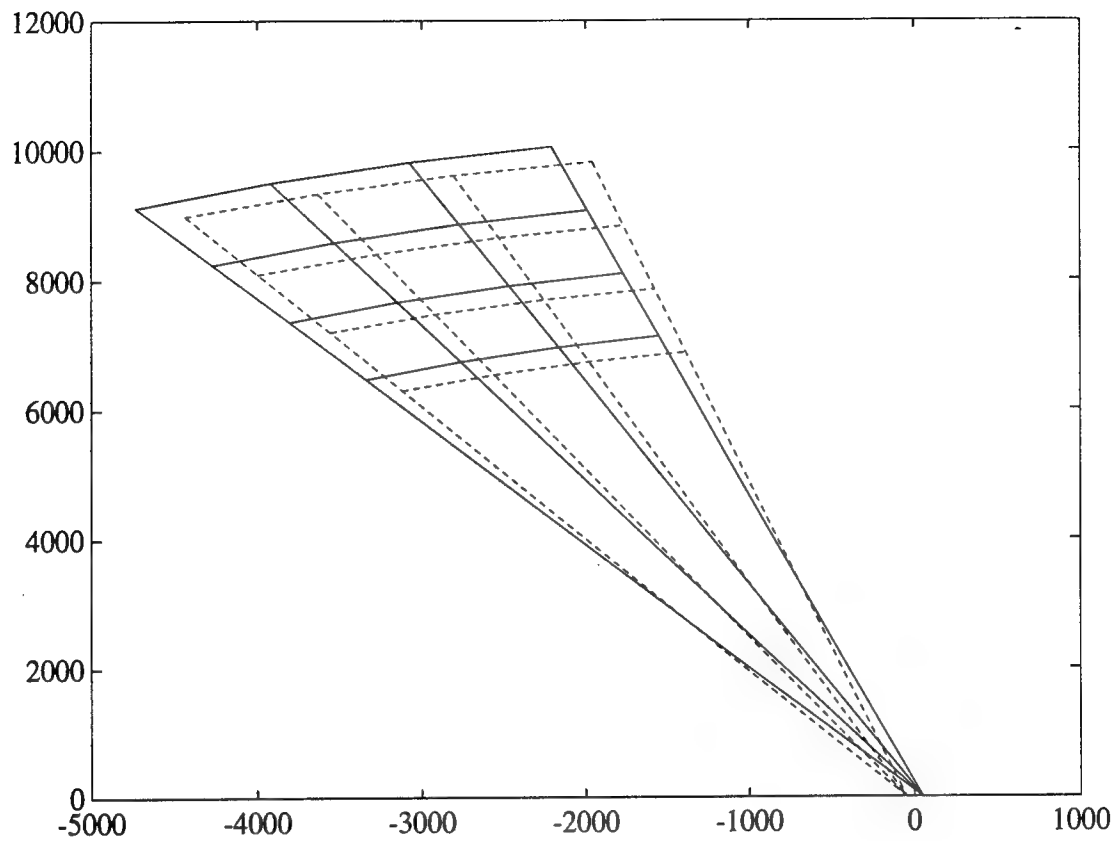
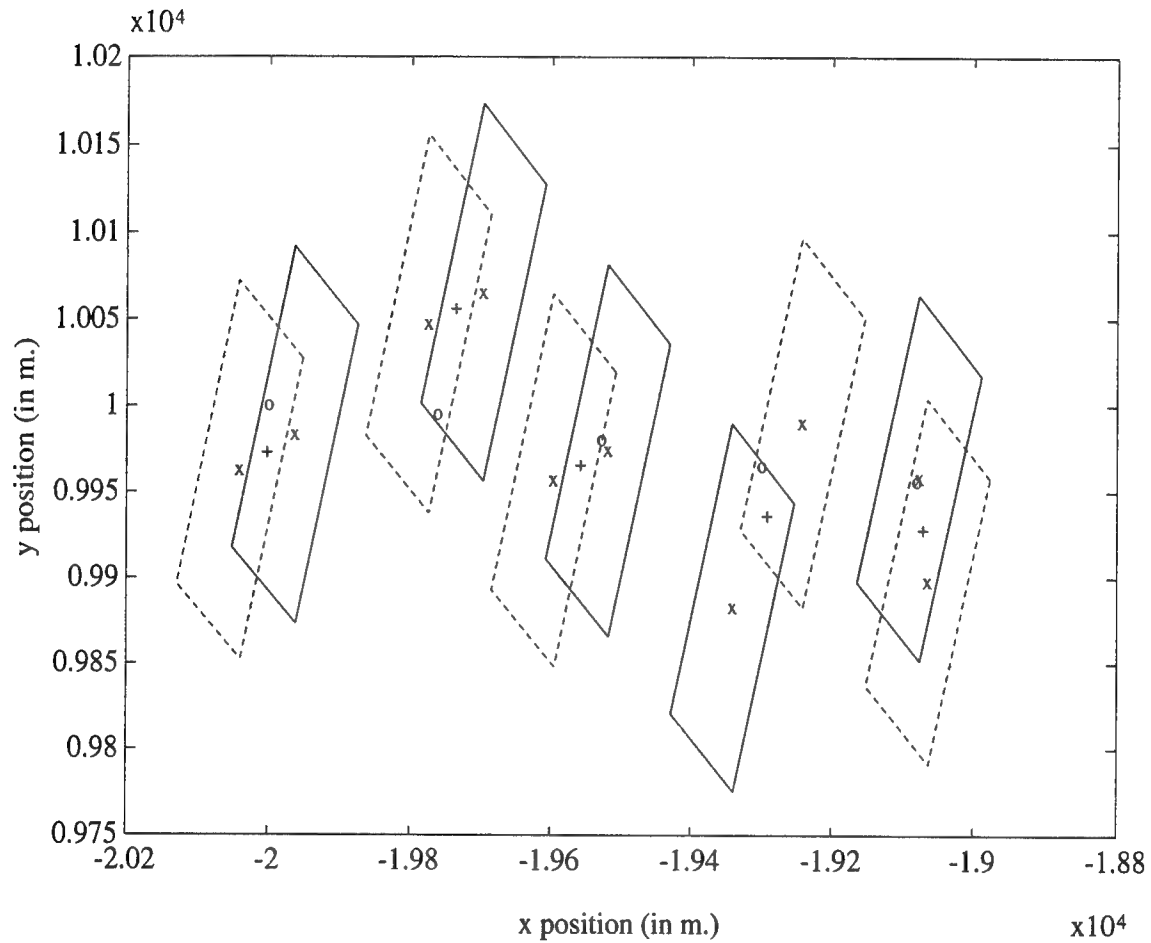Figure 1.1: Two close radars scanning a common (distant) area.

Figure 1.2: Resolution cells corresponding to two radars for an example track, along with the trajectory and measurements. ('o': the true trajectory, 'x': the centroids of resolution cells from individual measurements, '+': the average of both centroids or combined measurement.) The solid boxes represent hits from one radar, the dashed boxes hits from the other – note that the actual measurements lie within the overlaps of these.
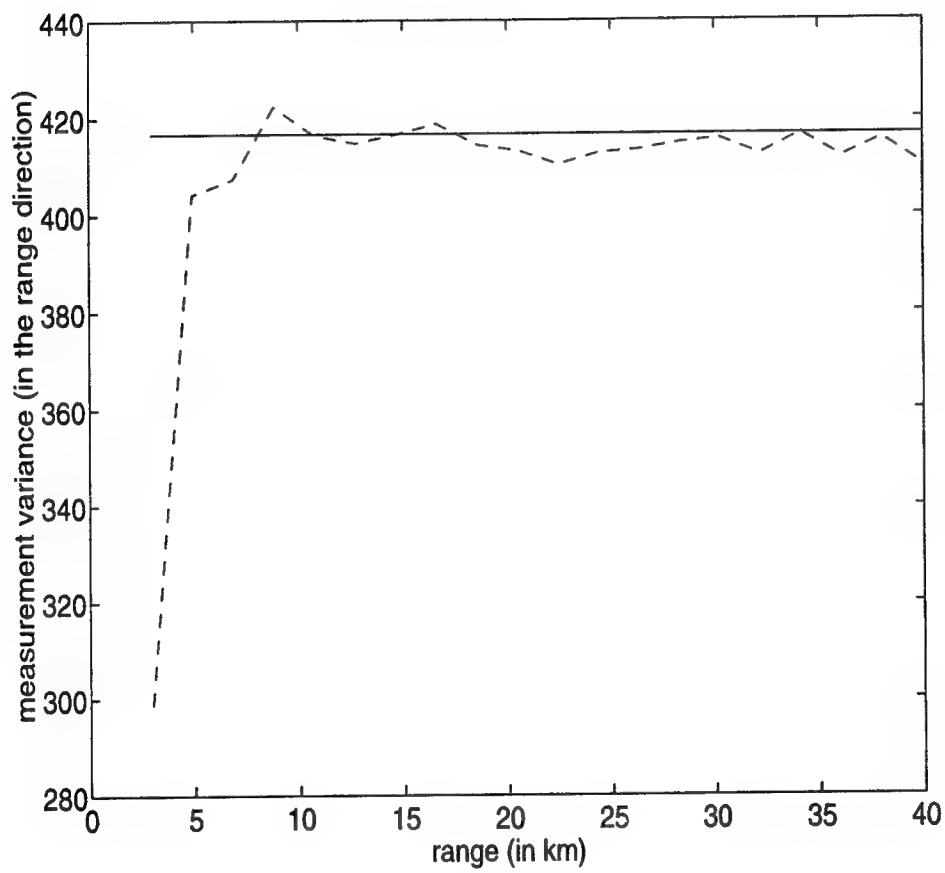
22

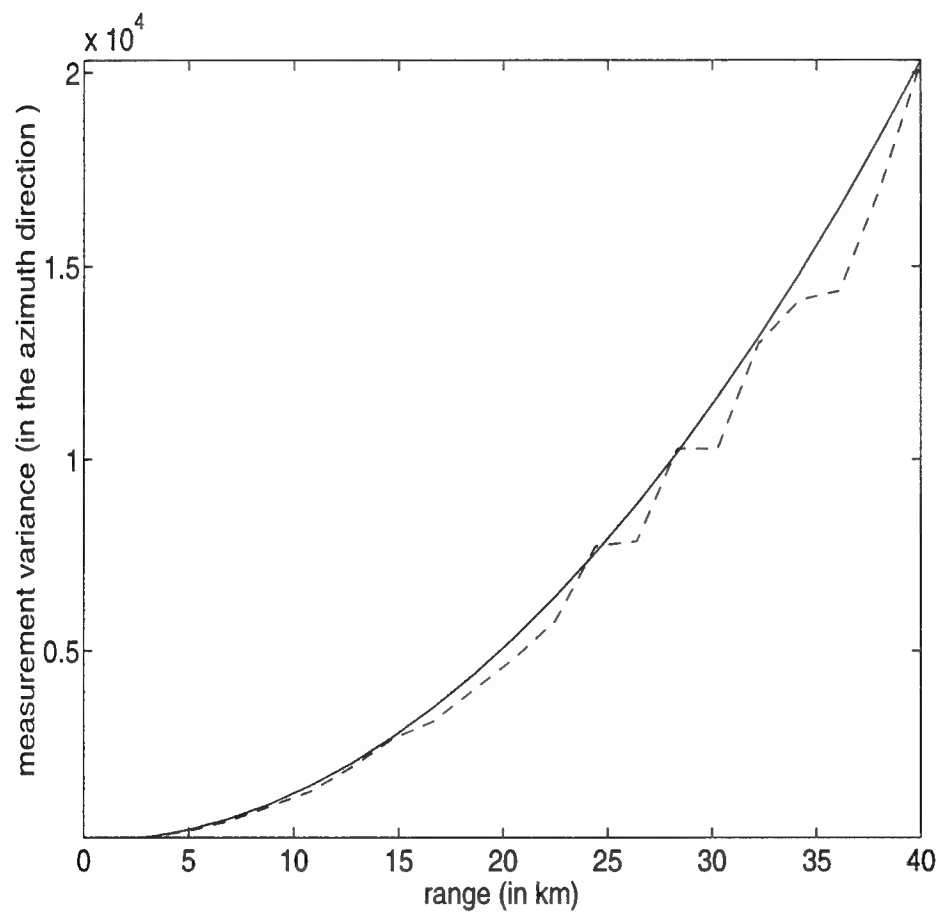Figure 1.3: Analytic (solid line) and computed (dashed line) measurement error variance in range.

Figure 1.4: Analytic (solid line) and computed (dashed line) measurement error variance in azimuth.
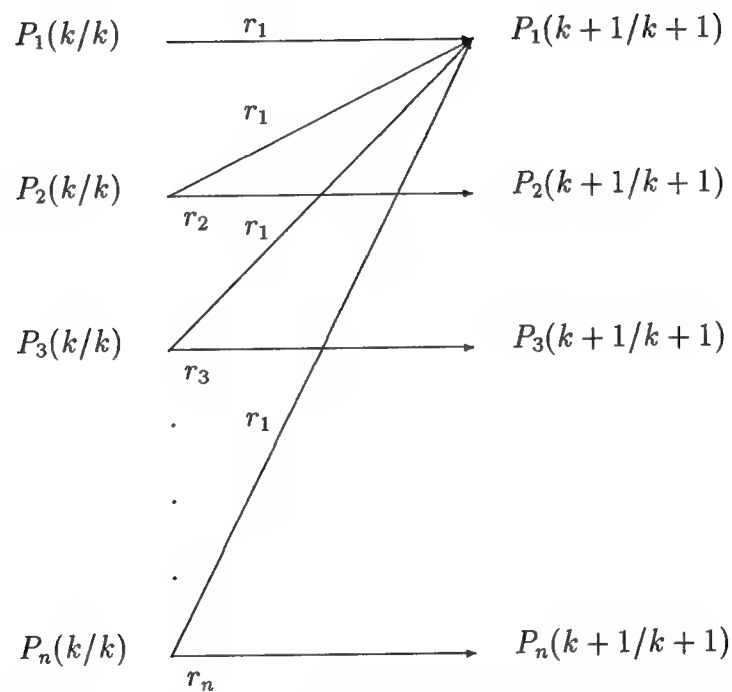
Figure 1.5: An illustration of the procedure for finding the steady-state covariance for a Kalman filter operating in a stochastic environment. Only selected transitions are shown.
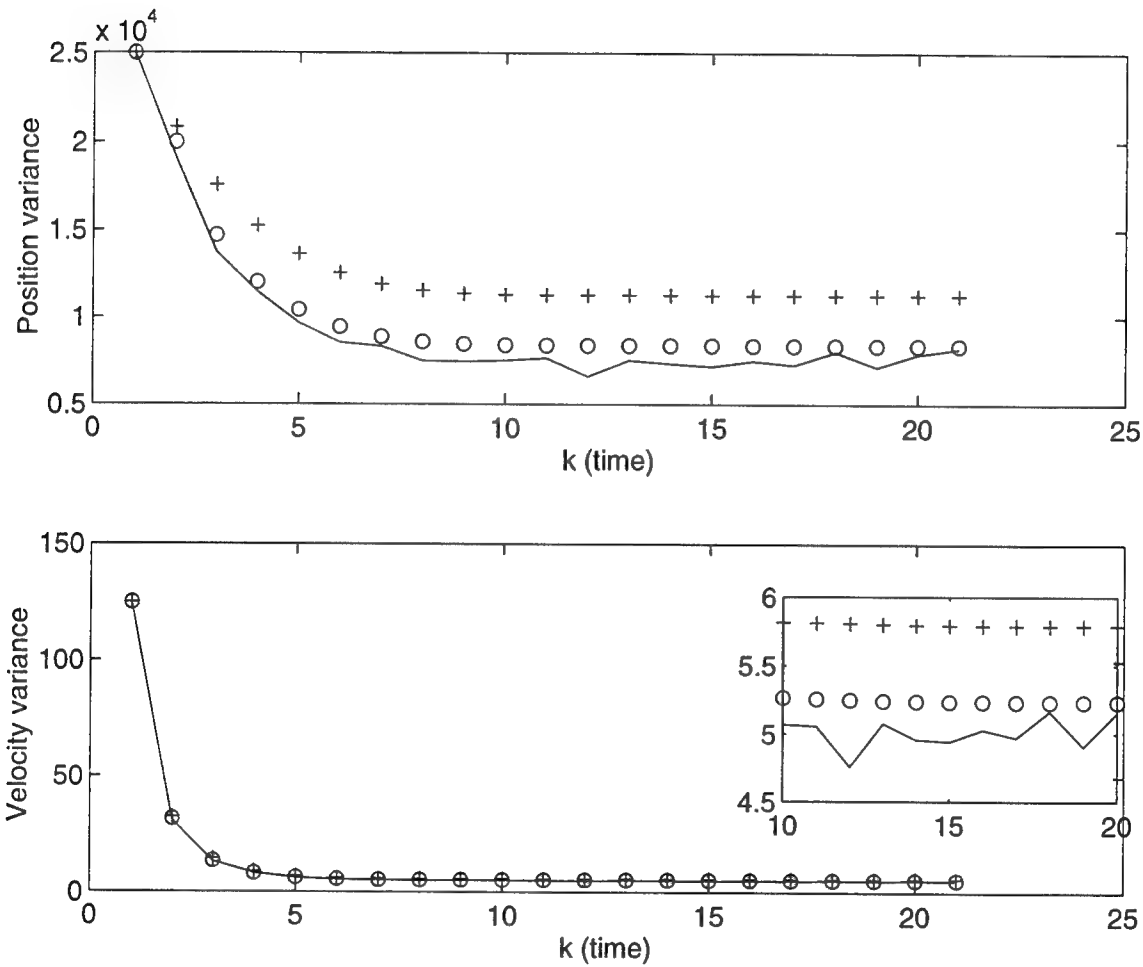
Figure 1.6: Evolution in time of the diagonal elements of $\mathbf{P}_{k|k}$. '+' is the steady state solution for $r = L^2/36$, 'o' is the approximation using 40 levels, r with density as in (1.4), '$--$' is the Monte Carlo simulation. The inset shows detail of the lower plot.

Figure 1.7: RMS position error as a function of the distance between radars (the range resolution is 100 meters and the azimuth resolution is 1 degree).
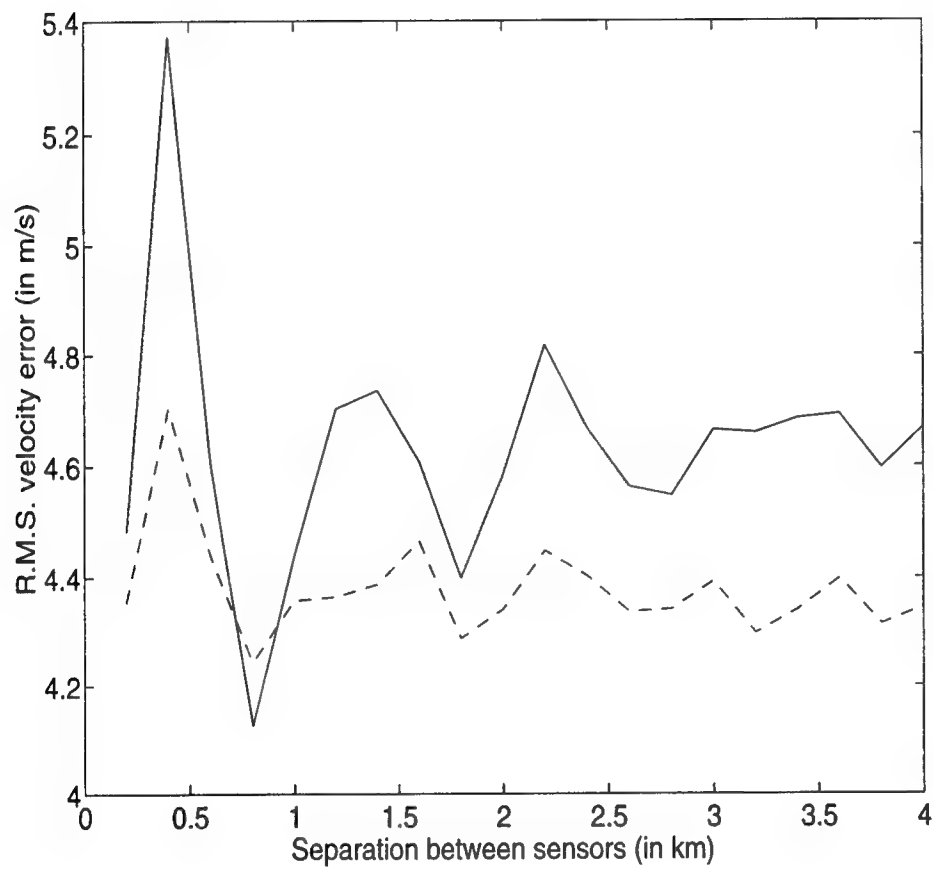
Figure 1.8: RMS velocity error as a function of the distance between radars (the range resolution is 100 meters and the azimuth resolution is 1 degree).

28

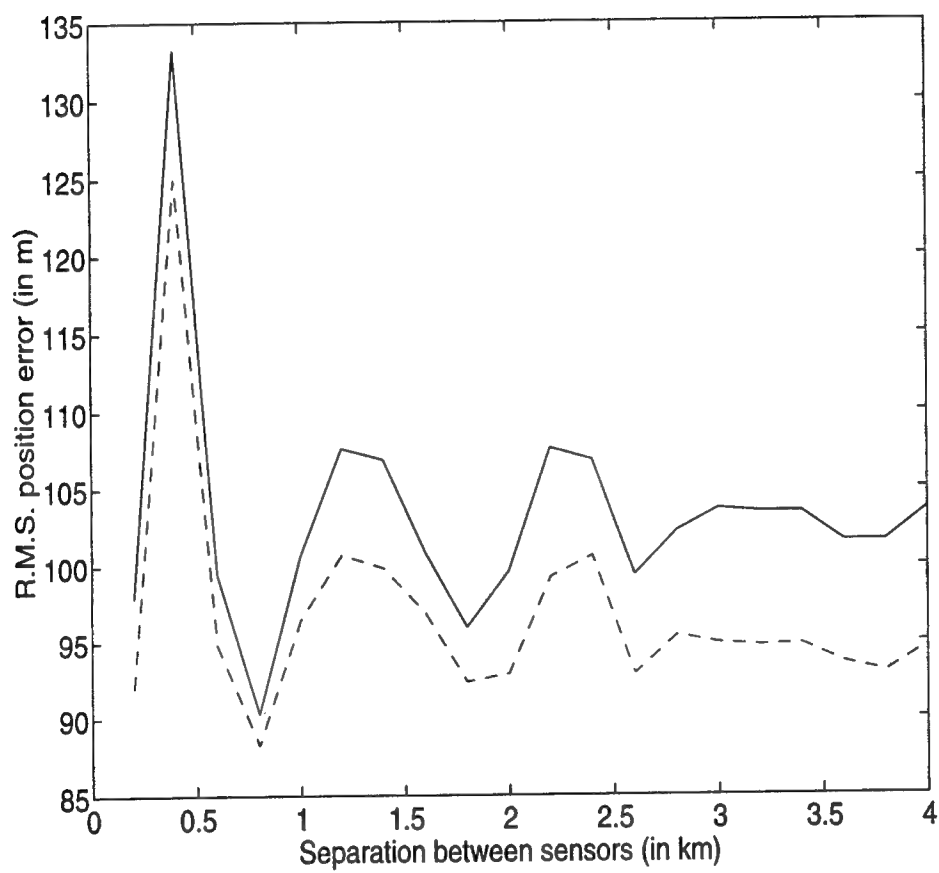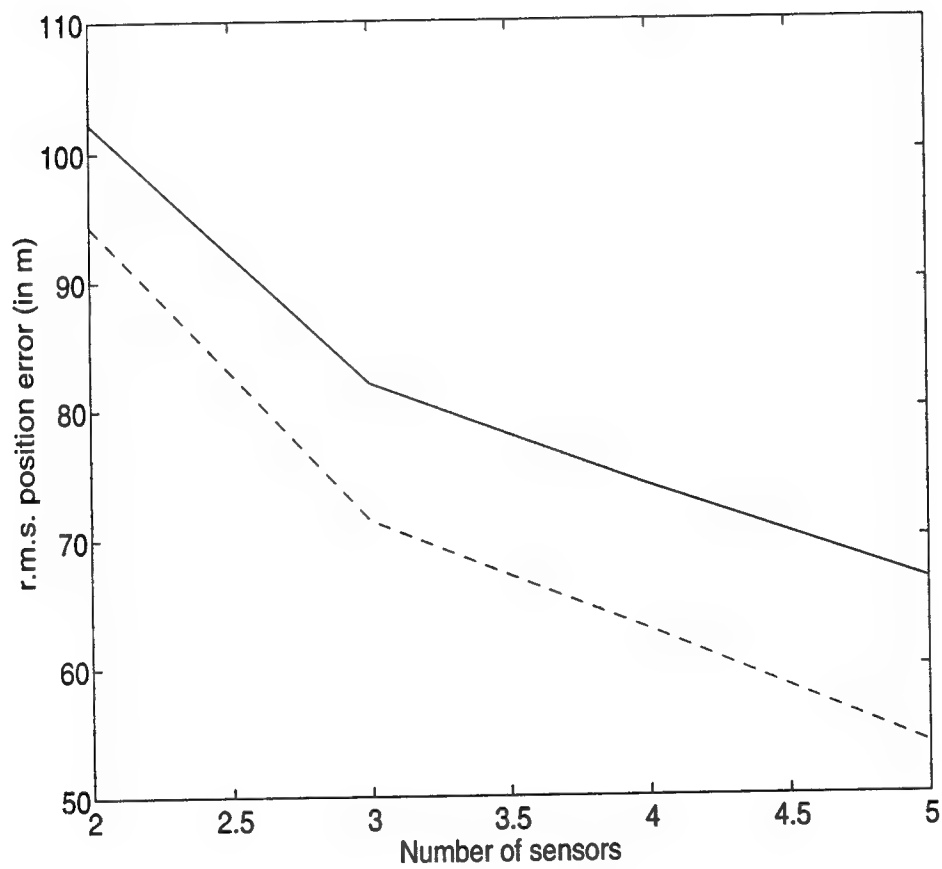Figure 1.9: RMS position error as a function of the number of radars (the distance between radars is 3.6 km, the range resolution is 100 meters, and the azimuth resolution is 1 degree).
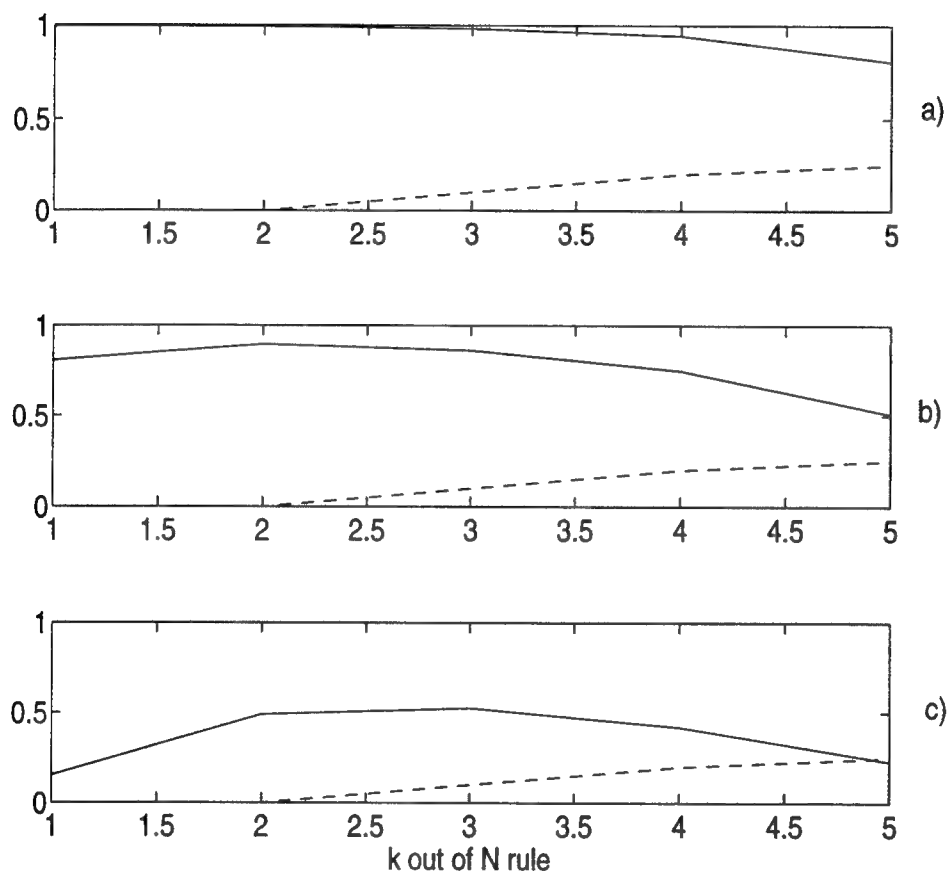
Figure 1.10: Probability of detection (solid line) and relative improvement in the measurement variance error (dashed line) for three different probabilities of false alarm: a) $P_F = .1$, b) $P_F = 1e - 3$, c) $P_F = 1e - 6$; SNR = 10 dB in the MBF case.
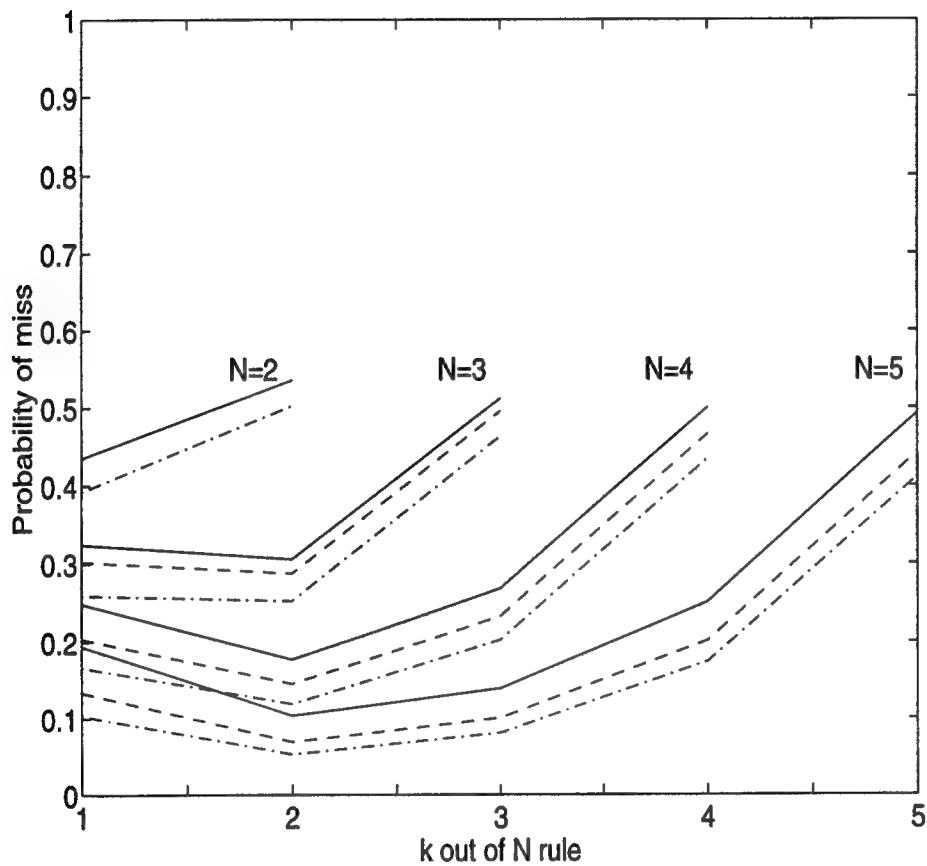
Figure 1.11: Probability of miss as a function of the k-out-of-N rule for 2, 3, 4 and 5 sensors (solid line: MBF fusion, dashed line: RCBF fusion, dashed-dot line: RCBF fusion taking into account the dynamic gain, SNR = 10 dB, Probability of false alarm = 1e-3). Note that for $N = 2$ MBF fusion and RCBF static fusion are identical.
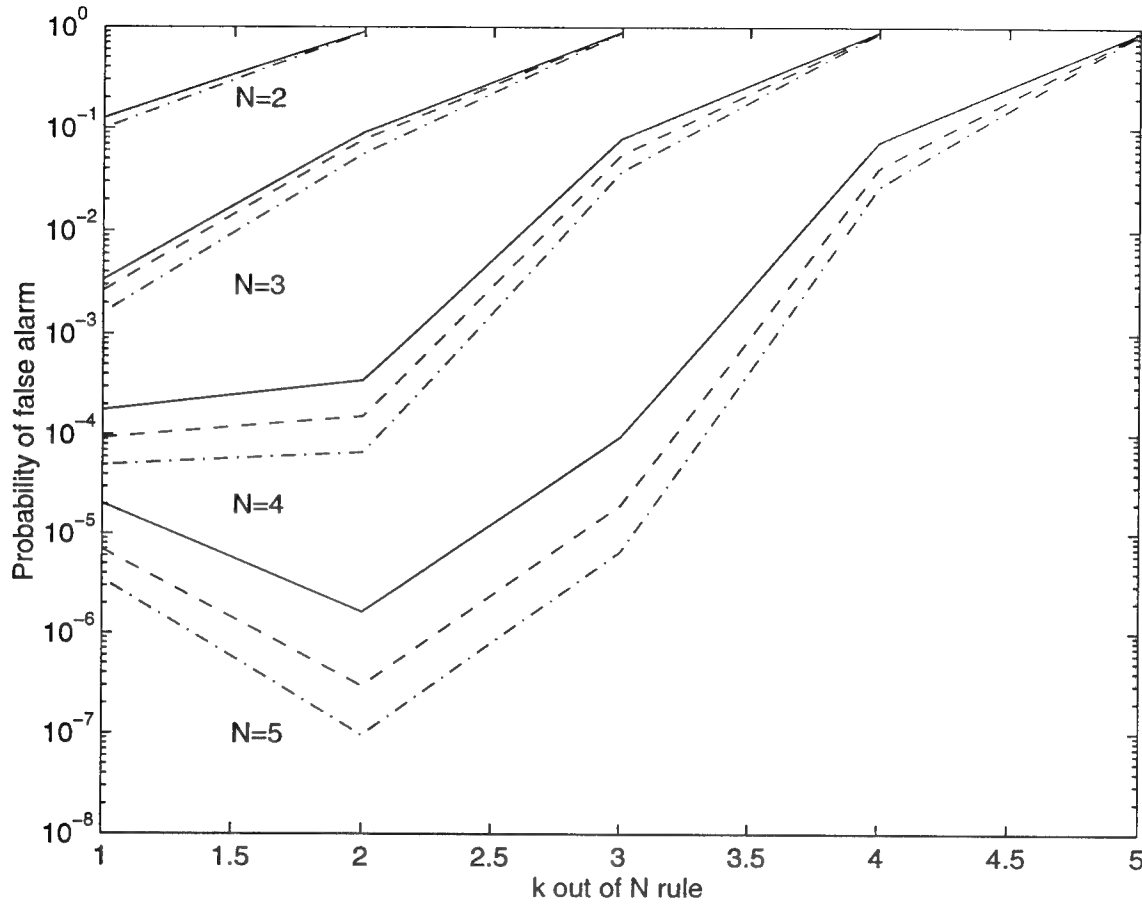
Figure 1.12: Probability of false alarm as a function of the k-out-of-N rule for 2, 3, 4 and 5 sensors (solid line: MBF fusion, dashed line: RCBF fusion, dashed-dot line: RCBF fusion taking into account the dynamic gain, SNR = 10 dB, Probability of detection = .999).
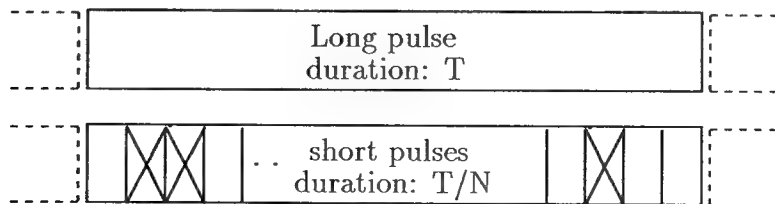
Figure 1.13: Long and short pulses to be used to illuminate the same region of the space consecutively.
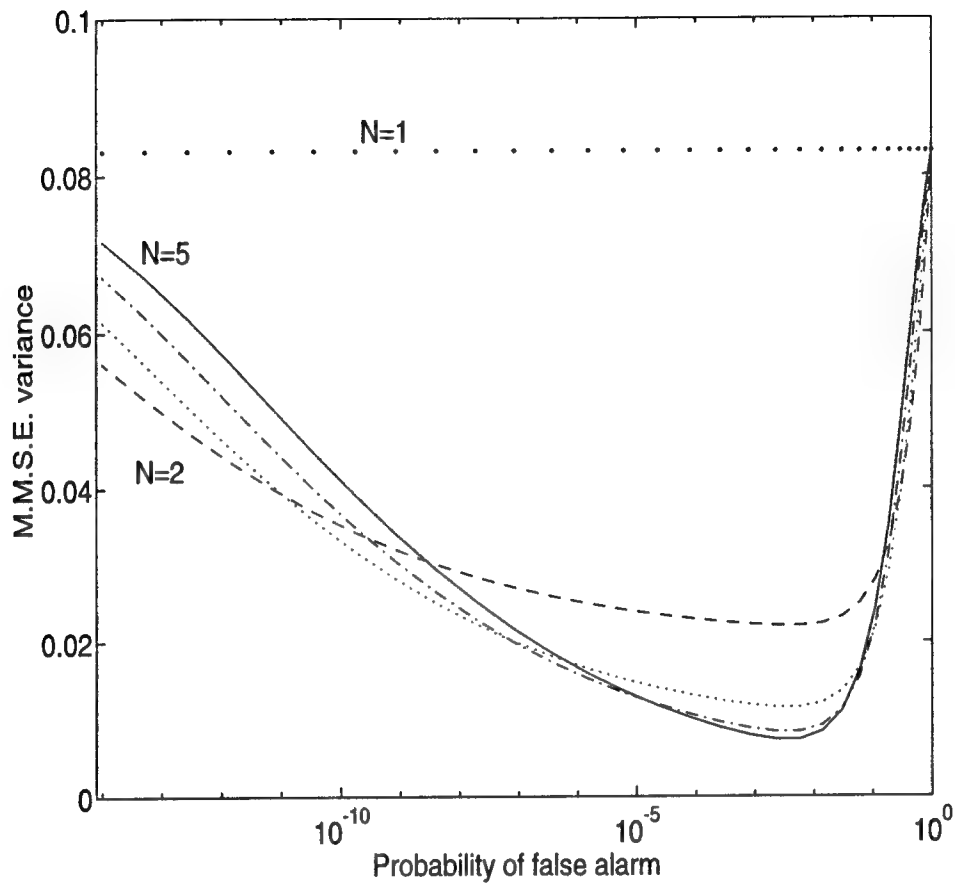
Figure 1.14: Variance of the range estimator *vs* the probability of false alarm for different relationships between the long and short pulse (N). The SNR for the long pulse is 30 dB
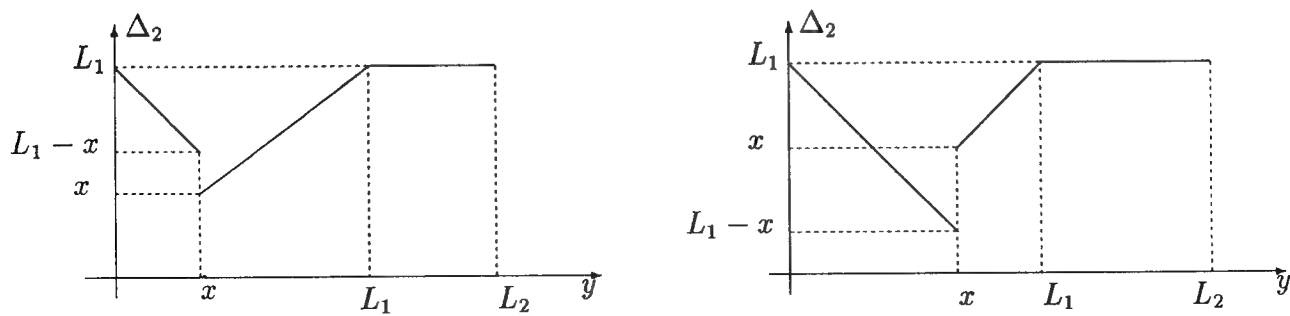
Figure 1.15: Length of the overlapping interval (Δ) as a function of the beginning of the second interval (y) for a fixed position of the target (x).

Chapter 2

The Coherent Case

Abstract

It is commonly understood that in active detection systems constant-frequency pulses correspond to good Doppler but poor delay resolution capability; and that linearly-swept frequency pulses have the opposite behavior. Many systems are capable of both types of operation, and hence in this report the *fusion* of such pulses is examined. It is discovered that in many (but not all) situations the features complement in such a way that tracking performance using a combined CW-FM pulse is improved by an order of magnitude when compared to a scheme using only a full CW or FM pulse. Also investigated are alternating-pulse systems, and while these are suboptimal their performances appear robust.

2.1   Introduction

Many active detection systems (radar and sonar) are capable of waveform agility – particularly of constant frequency (CW) and linearly-swept frequency (FM, or chirp) – in their transmitted pulses. It appears that such systems presently use their abilities, but do not *fuse* the results in an optimal way. The purpose of this report is to examine the possible benefits from such fusion. To be specific, we shall assume:

- A Swerling I target and white Gaussian noise. A target made up of many reflectors simplifies our analysis, as will be seen. We expect that results using other models would give rise to similar conclusions.

- A combined CW-FM transmitted pulse. A system implementation may be such that CW and FM pulses are alternately transmitted, but it simplifies our intuition to assume that pulses are abutted. We do not consider phase-coded or other broadband pulses.

- Perfect correlation between the target RCS for the CW and FM pulses. In the case of our "abutted" pulses this makes sense, as one would expect little change in the observed target scattering characteristics during the course of a pulse.

- Single-pulse operation. We do not investigate the effect of multiple-pulse integration, as it is a complication and we do not expect qualitative differences in our results.

- Known ambient noise power. Incorporating CFAR processing adds very little to our results.

The apportionment of energy between CW and FM sub-pulses will be a parameter, and as such we include full-CW and full-FM as special cases.

It is "common knowledge" that a constant-frequency (CW) pulse has good range-rate (i.e. Doppler frequency) resolution characteristics, but is relatively poor at range resolution; conversely, that an FM pulse can have excellent range, but comparatively poor range-rate resolutions. The idea of this research is to investigate whether a *combined* system using

35

both kinds of pulses – and consequently taking advantage of their complementary features – can perform better than either extreme.

Our goal, therefore, is to compare the performances, with respect to tracking error, of a number of schemes:

- A fully CW pulse, of duration T_p.

- A fully linear-FM pulse, of duration T_p; sweep rate is a parameter.

- A combined FM-CW pulse, with the first $(1 - \kappa)T_p$ seconds being of constant frequency, and the remaining κT_p of linearly-swept frequency. As above, sweep rate is a parameter.

- FM and CW pulses, each of duration T_p, alternating on successive scans.

Comparison will be on the basis of tracking performance as specified by the algebraic Riccati equation (ARE) modified by the *hybrid conditional averaging* (HYCA) technique to account for missed detections.

In section 2.2 we give mathematical preliminaries and signal models; naturally, we discuss ambiguity functions and probability of detection. Section 2.3 details our analysis procedure, specifically relating to our assumptions about measurement accuracy and the use of the HYCA approach. Finally, section 2.4 gives our results and recommendations.

2.2   Background

2.2.1   Signal and Receiver Model

Receiver processing is assumed, without loss of generality, to be performed at "baseband"; that is, subsequent to carrier-removal via mixing. With $r_i(t)$ the received waveform, we have, as an output process of our matched filter,

$$
\begin{aligned}
x(t) &= \int s(t - u)^* \left(r_i(u) e^{-j2\pi f_c u} \right) du \\
&\equiv \int s(t - u)^* r(u) du
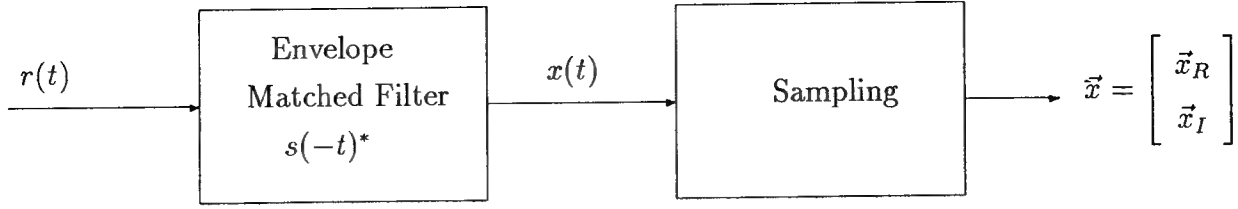\end{aligned}
\tag{2.1}
$$

36

Figure 2.1: Complex envelope matched filter operation and sampling

where f_c is the carrier frequency, and $s(t)$ is a replica of the transmitted baseband signal. This standard scheme is as pictured in figure 2.1.

The (baseband) received signal will be modeled as a return from a Swerling I [4] target;

$$
\begin{aligned}
r(t) &= \left(\sum_i A_i \right) s(t, \tau, f_d) \mathcal{I}(\text{target}) + \nu(t) \\
&= A\, s(t - \tau) e^{j2\pi f_d t} \mathcal{I}(\text{target}) + \nu(t)
\end{aligned}
\tag{2.2}
$$

where $s(t, \tau, f_d) = s(t - \tau) \exp(j2\pi f_d t)$ is a delayed and Doppler-shifted replica of the emitted baseband complex envelope signal $s(t)$, τ is the (unknown) delay and f_d is the (unknown) Doppler frequency shift; $\mathcal{I}(\text{target})$ is a target indicator. The complex numbers A_i represent the amplitude and phase of each of the target reflectors; for a Swerling I model the assumption is that there are many such reflectors, each one with a random amplitude and phase. Hence $A = \sum_i A_i$ approaches a complex Gaussian random variable with zero mean and variance $2\sigma_A^2$; $\nu(t)$ is complex white Gaussian noise independent of A, with zero mean and variance $2N_o$. The signal $s(t)$ comprises a CW pulse (of duration $(1 - \kappa)T_p$) followed by an FM pulse (of duration κT_p), and phase-continuity is assumed; in the case of a purely-FM or purely-CW pulse, we take respectively $\kappa = 0$ or $\kappa = 1$.

2.2.2   Ambiguity Function

Of interest to us in the sequel is the "ambiguity function" [10], given by

$$
\mathcal{A}(\tau, f_d) = \left| \int s(t) s^*(t - \tau) e^{j2\pi f_d t} dt \right|^2
\tag{2.3}
$$

As is clear, the ambiguity function specifies the output of the matched filter in the absence of noise; its shape is related to the implied resolution cell, and is strongly dependent on the

waveform. Examples showing the effect of different energy apportionments between CW and FM pulses are given in figures 2.3, 2.4, 2.5, and 2.6.

2.2.3 Detection

Here we discuss the detection of a delayed and Doppler-shifted return. The implied test will use the output of the envelope (magnitude square of the real and imaginary parts according to figure 2.1) matched filter. At time t the magnitude is

$$|x(t)|^2 = \left| \int_0^t r(\lambda)s^*(\lambda - (t - T_p))d\lambda \right|^2 \tag{2.4}$$

Since the delay and Doppler shift are unknown, we examine without loss of generality this at time T_p, and hence the test is given by

$$|x(T_p)|^2 = \left| \int_0^{T_p} r(\lambda)s^*(\lambda)d\lambda \right|^2 \mathop{\lessgtr}_{K_1}^{H_0} \mathcal{T} \tag{2.5}$$

with $\mathcal{T}$ a threshold. We will characterize now the random variable $x(T_p)$ under the noise-only hypothesis (H) and under the target-present hypothesis (K).

Noise-only hypothesis. Here we have, according to our formalism, $\mathcal{I}(target) = 0$, and hence

$$x(T_p) = \int_0^{T_p} \nu(\lambda)s^*(\lambda)d\lambda \tag{2.6}$$

The random variable $x(T_p)$ is complex Gaussian, with zero mean and variance given by:

$$
\begin{aligned}
\sigma_0^2 = E\left\{x(T_p)x^*(T_p)\right\} &= \int_0^{T_p}\int_0^{T_p} E\left\{\nu(\lambda)s^*(\lambda)\nu^*(\alpha)s(\alpha)\right\} d\lambda d\alpha \\
&= 2N_o \int_0^{T_p} |s(\lambda)|^2 d\lambda \\
&= 2N_o\xi
\end{aligned}
\tag{2.7}
$$

The last integral, ξ, is the energy of the transmitted pulse.

Target Present. In this case we have $\mathcal{I}(target) = 1$, and the matched filter output is thus

$$x(T_p) = \int_0^{T_p} \left[As(\lambda - \tau)e^{j2\pi f_d\lambda} + \nu(\lambda) \right] s^*(\lambda)d\lambda \tag{2.8}$$

This random variable is still zero mean, with variance given by:

$$
\begin{aligned}
\sigma_1^2 = E\left\{x(T_p)x^*(T_p)\right\} &= \int_0^{T_p}\int_0^{T_p} s^*(\lambda)s(\alpha) \\
&\quad E\left\{[As(\lambda-\tau)e^{j2\pi f_d\lambda}+\nu(\lambda)][As(\alpha-\tau)e^{j2\pi f_d\alpha}+\nu(\alpha)]^*\right\}d\lambda d\alpha \\
&= 2N_o\int_0^{T_p}|s(\lambda)|^2 d\lambda + 2\sigma_A^2\left|\int_0^{T_p} s^*(\lambda)s(\lambda-\tau)e^{j2\pi f_d\lambda}d\lambda\right|^2 \\
&= 2N_o\xi + 2\sigma_A^2\xi^2\mathcal{A}(\tau,f_d) \\
&= \sigma_0^2(1 + \frac{2\sigma_A^2\xi}{\sigma_0^2}\mathcal{A}(\tau,f_d))
\end{aligned}
\tag{2.9}
$$

where we have used the definition of the ambiguity function $\mathcal{A}$ given by (2.3) (with a normalization factor to take into account the non-unity energy signals).

Recall that the magnitude square of a complex Gaussian random variable is exponentially distributed, with density given by:

$$
x = \mathcal{N}(0,\sigma_i^2) \rightarrow \quad y = x^2 \sim \frac{1}{2\sigma_i^2}e^{-y/2\sigma_i^2}
\tag{2.10}
$$

We consequently have

$$
P_{fa} = \int_T^\infty \frac{1}{2\sigma_0^2}e^{-x/2\sigma_0^2}dx = e^{-T/2\sigma_0^2}
\tag{2.11}
$$

and

$$
P_d = \int_T^\infty \frac{1}{2\sigma_1^2}e^{-x/2\sigma_1^2}dx = e^{-T/\left(2\sigma_0^2(1+\frac{2\sigma_A^2\xi^2}{\sigma_0^2}\mathcal{A}(\tau,f_d))\right)}
\tag{2.12}
$$

using (2.11) we get

$$
P_d = P_{fa}^{1/(1+\frac{2\sigma_A^2\xi^2}{2N_0\xi}\mathcal{A}(\tau,f_d))} = P_{fa}^{1/(1+\text{SNR }\mathcal{A}(\tau,f_d))}
\tag{2.13}
$$

This relationship has been plotted for a particular probability of false alarm, as a function of τ and f_d and for different CW-FM apportionments in figures 2.7, 2.8, 2.9, and 2.10. Observe the pronounced variation in the shape of the "resolution cells" as a function of the apportionment.

2.3 Comparison of Schemes

In this section we give details relating to our mode of comparison of our various pulse-fusion schemes. Of clear interest are the signal model, the target model, the measurement model, and the analytical framework; and we devote a subsection to each.

In our studies we use a baseband signal $s(t)$ comprising $(1 - \kappa)T_p$ seconds of "CW" (which at baseband is constant), followed by κT_p seconds of linearly-swept FM. That is, we have

$$
s(t) = \begin{cases}
\frac{1}{\sqrt{T_p}} & 0 \leq t < (1 - \kappa)T_p \\
\frac{1}{\sqrt{T_p}} e^{j\left(\pi k \left(t - \frac{(2-\kappa)T_p}{2}\right)^2 + \theta\right)} & (1 - \kappa)T_p \leq t < T_p \\
0 & else
\end{cases}
\tag{2.14}
$$

where θ is chosen to ensure phase-continuity. Results using other configurations, such as FM followed by CW, exhibit different ambiguity functions; however, results are qualitatively similar.

2.3.1 The Target Model

Our target model is kinematic with range/range-rate as state variable. That is, we have:

$$
\mathbf{x}(t + 1) = \mathbf{F}\mathbf{x}(t) + \mathbf{G}\mathbf{u}(t) + \mathbf{v}(t)
$$

where the process noise is white and zero-mean with autocorrelation $E\{\mathbf{v}(t)\mathbf{v}^T(t)\} = \mathbf{Q}(t)$. The "ownship" motion $\mathbf{u}(t)$ is taken as zero without loss of generality. Our observations are delay and Doppler-shift:

$$
\mathbf{y}(t) = \mathbf{H}\mathbf{x}(t) + \mathbf{w}(t)
\tag{2.15}
$$

where the measurement covariances $E\{\mathbf{w}(t)(\mathbf{w}(t))^T\} = \mathbf{R}$ are functions of the pulse shape employed. We write

$$
\mathbf{F} = \begin{bmatrix} 1 & \Delta t \\ 0 & 1 \end{bmatrix}
$$

$$\mathbf{H} = \begin{bmatrix} 2/c_o & 0 \\ 0 & -2f_o/c_o \end{bmatrix}$$

$$\mathbf{Q} = \sigma_q^2 \begin{bmatrix} \frac{\Delta t^3}{3} & \frac{\Delta t^2}{2} \\ \frac{\Delta t^2}{2} & \Delta t \end{bmatrix} \tag{2.16}$$

where Δt represents the time between samples, c_o is the wave's propagation speed in the medium, and f_o is the carrier frequency. Note that for clarity we have chosen to work in one dimension, range, since it is in this domain that the differences between the various waveform types become apparent.

2.3.2 The Measurement Model

In the most generous world it would be possible to determine an optimal (presumably maximum-likelihood) estimate of delay τ and shift f_d for each detection. However, due to the complexity of such an operation most systems resort to a sampled or quantized approach, which immediately gives rise to the concept of a *resolution cell*. For us, a resolution cell defines a region in range/range-rate space which can be tested for the presence or absence of a target. If a detection is declared, then the corresponding measurement as supplied to the tracking algorithm is the *centroid* of that resolution cell, with measurement covariance $\mathbf{R}$ derivable from the cell's size and shape, assuming a uniform distribution of the measurements along the cell. Ideally, therefore, resolution cells should be sufficiently regular in geometry that they are mutually-exclusive and exhaustive of space (i.e. they "tesselate"); and that a target located in a given cell will not effect a detection in any other. Unfortunately this is overly-idealized even in the extreme CW or FM cases, in which constant-P_d contours tend to be elliptical in shape. And it is certainly not true when a combined pulse is used.

For the current research we have adopted the policy of assuming a resolution cell to be defined by all delay- and Doppler-space for which P_d exceeds a certain threshold. It should further be noted that when this threshold is chosen too low resolution cells can became bizarrely-formed, and may even fail to be simply-connected. As such, we concentrate on cases in which this does not happen.

41

Our steps are:

1. Define a resolution cell by selecting a "detection threshold" within which the probability of detection is consistently exceeded.

2. For such a resolution cell, define the probability of detection to be its average value within that area.

3. For such a resolution cell, define the measurement covariance by assuming the detection-producing target's location to be uniformly-distributed within that area.

Our resolution cells do not in general tesselate; however, realistically, neither do those of the simple CW or FM schemes.

2.3.3 Analytical Approach

The state estimation-error covariance matrix is updated by the Kalman filter as follows ([11]):

$$\mathbf{P}(t+1|t) = \mathbf{F}\mathbf{P}(t|t)\mathbf{F}' + \mathbf{Q}$$
$$\mathbf{S}(t+1) = \mathbf{H}\mathbf{P}(t+1|t)\mathbf{H}' + \mathbf{R}(t)$$
$$\mathbf{W}(t+1) = \mathbf{P}(t+1|t)\mathbf{H}'\mathbf{S}(t+1)^{-1}$$
$$\mathbf{P}(t+1|t+1) = \mathbf{P}(t+1|t) - \mathbf{W}(t+1)\mathbf{S}(t+1)\mathbf{W}'(t+1) \tag{2.17}$$

With $\mathbf{R}(t) = \mathbf{R}$ this evolves to a steady state which characterizes the performance of the filter, and this solution can be computed from the associated algebraic Riccati equation (ARE). However, in the case that detections can be missed the measurement noise has a stochastic covariance ($\mathbf{R}(t)$), and there *is* no steady-state. However, it is legitimate to seek the *expected value of the error-covariance*.

Let us assume that the measurement noise covariance matrix can be modeled as:

$$\mathbf{R}(t) = \begin{cases} \mathbf{R} & detection \\ \infty\mathbf{I} & miss \end{cases} \tag{2.18}$$

42

where, as described in the previous subsection, $\mathbf{R}$ is derived from the size and shape of the resolution cell and the probabilities of detection and miss are the corresponding averaged values.

The method here employed is the HYCA (hybrid conditional average) method of [8]. Under this scheme we embed a binary random process $\{s(t)\}$ into the Riccati equation such that

$$
s(t) \;=\; \left\{ \begin{array}{ll} 1 & detection \\ 2 & miss \end{array} \right. \qquad Pr(s(t) = i) \;=\; \left\{ \begin{array}{ll} P_d & i = 1 \\ 1 - P_d & i = 2 \end{array} \right.
$$
$$(2.19)$$

Then, with $\mathbf{P}_{ij}(t+1|t+1)$ defined as the estimation error covariance given $s(t) = i$ and $s(t+1) = j$ (obtainable from (2.17) with $\mathbf{R}(t)$ according to (2.18)), we get

$$
\mathbf{P}_j(t+1|t+1) = \sum_{i=1}^{2} \mathbf{P}_{ij}(t|t) Pr(s(t) = i) \qquad (2.20)
$$

The matrices $\mathbf{P}_i(t|t)$ reflect the state error covariance given that $s(t) = i$. The mean state covariance matrix (*i.e.* what we want) is given by

$$
\mathbf{P}(t|t) = \sum_{j=1}^{M} \mathbf{P}_j(t|t) Pr(s(t) = i) \qquad (2.21)
$$

evaluable when this matrix Markov process reaches steady state. Figure 2.2 shows a flow-diagram that reflects the mechanics of the iterations.

2.4 Results

We have developed two sets of representative plots: one with parameters suitable for sonar (i.e. low carrier frequency, long pulses, slow-moving targets), the other more applicable to the radar situation (high frequency, short pulses, fast-moving targets).

2.4.1 The Sonar Case

For the active sonar case we assume a pulse-length of $T_p = 1s$, a sampling rate of $\Delta t = 30s$, and a carrier frequency of $f_d = 3.5\text{kHz}$. For various values of κ we show in figures 2.7, 2.8,
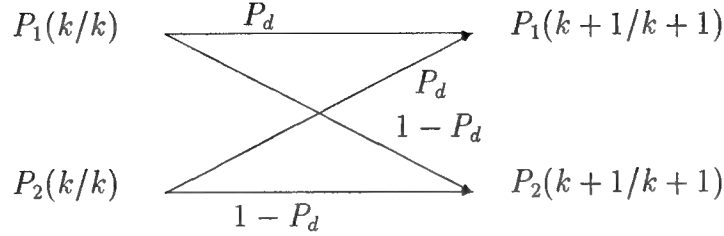
43

Figure 2.2: An illustration of the procedure for finding the steady-state covariance for a Kalman filter operating in a stochastic environment.

2.9, and 2.10 representative constant-P_d contours; in each case here we have SNR= 20 dB and $k = 50$Hz/s – this corresponds to a maximum frequency sweep of 50Hz, found in case $\kappa = 1$. As discussed earlier, the "resolution cell" is determined by the volume enclosed by one of these contours, and for the CW and FM cases these are as expected: in the former we have good and comparatively poor range resolutions; and in the latter apparently the reverse. Intermediate values of κ give more complicated P_d contours, but if a large contour is chosen as the resolution cell boundary the results appear promising. Here the resolution cell is determined by the $P_d = 0.8$ contour. In figure 2.11 we show the corresponding measurement errors – these are the diagonal elements of the measurement covariances $\mathbf{R}$.

In figure 2.12 we show steady-state tracking errors obtained via our HYCA procedure, for various values of maneuvering index (i.e. σ_q^2). There are a number of items of note here:

- For a low maneuvering index, corresponding to relatively straight-line target motions, FM outperforms CW. However, for more maneuvering targets the opposite is true. This perhaps-surprising behavior stems from shape of the FM resolution cell as shown in figures 2.7, 2.8, 2.9, and 2.10: in the former case where tracking is accurate the measurement range uncertainty corresponds to a "strip" across a narrow resolution cell; in the latter the prior tracking information contributes little to knowledge of position within the resolution cell, whose elongated shape in turn produces an

44

inaccurate measurement.

- In all cases shown here a value $\kappa = 50\%$ (half FM, half CW) is optimal. The improvement in range estimation, as compared to FM-only or CW-only can be as high as *an order of magnitude.*

- The alternating-pulse scheme, pictured as the horizontal lines in figure 2.12, has excellent performance for low-maneuverability targets.

Figure 2.13 corresponds to figure 2.12, with the difference that the resolution here is determined by the $P_d = 0.6$ contour. In this case we find that a combined pulse, far from contributing an order-of-magnitude reduction to the estimation errors, actually causes a degradation in performance for low-maneuvering targets. The reason for this is clear from figures 2.7, 2.8, 2.9, and 2.10: for lower values of P_d the resolution cell structure becomes complicated and no longer simply connected, and the result is a high level of measurement uncertainty from a given detection.

2.4.2   The Radar Case

In this case we have assumed a rather long pulse, $T_p = 100\mu s$ and a carrier frequency of $f_d = 40\text{GHz}$. The long pulse combined with the high carrier frequency gives a frequency resolution well below the expected Doppler-shift. The FM modulation corresponds in this case to a $k = 500\text{kHz/s}$. The sampling period is $\Delta t = 1s$, and the SNR, as before is 20 dB. Figure 2.14 shows the steady state prediction variances for three different process noise intensities. As we can see, for low process noise, the alternating scheme has a performance almost as good as the combined pulse. On the other hand, for high process noise, the combined pulse performs considerable better than the alternating scheme.

2.5 Summary

In this report we have investigated the benefits from combined constant- and swept-frequency pulse fusion systems. Our comparison is in terms of steady-state tracking performance, via the algebraic Riccati equation, as modified to account for missed detections. Fundamental to our analysis is the resolution cell – defined for our purposes as the volume of range/range-rate space wherein the probability of detection exceeds a given threshold – and its implied measurement error. Our results indicate that in certain reasonable situations a fused pulse can lead to dramatic improvement in tracking error – as much as an order of magnitude range-estimation error variance reduction.
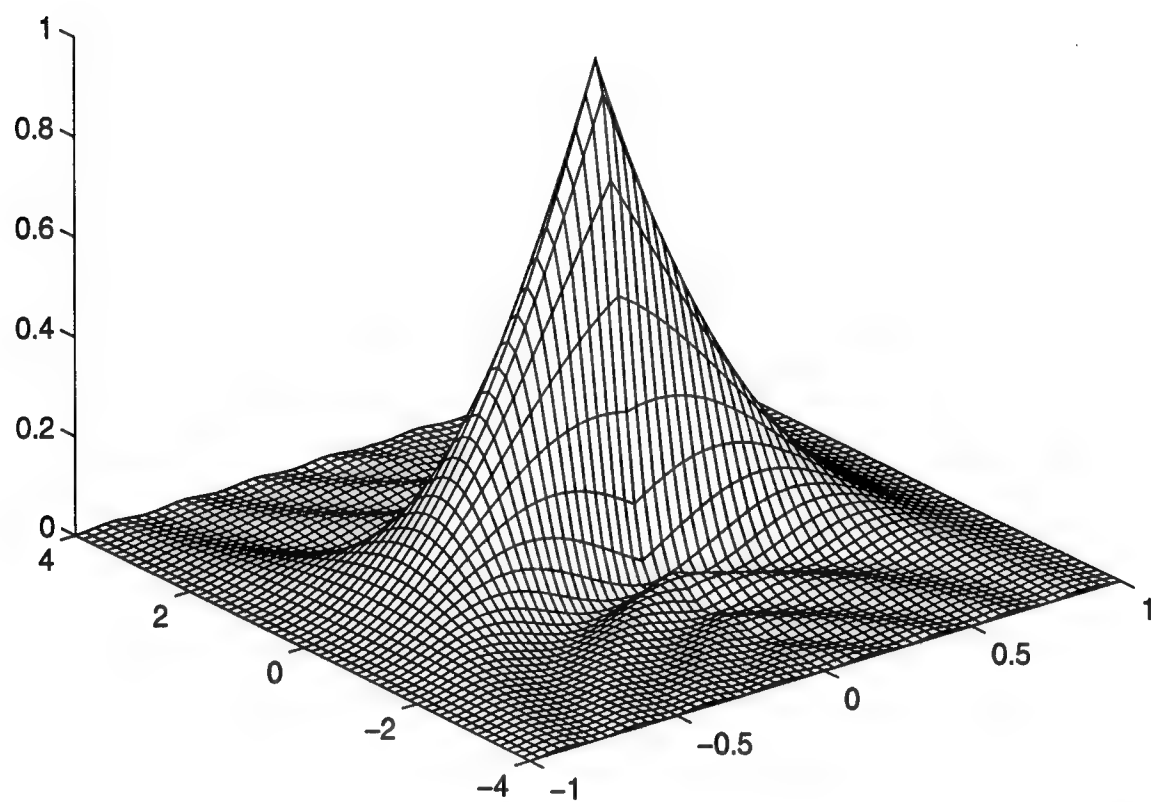
Figure 2.3: Ambiguity function for frequency sweep rate $k = 50\text{Hz/s}$: $\kappa = 0\%$ (CW).
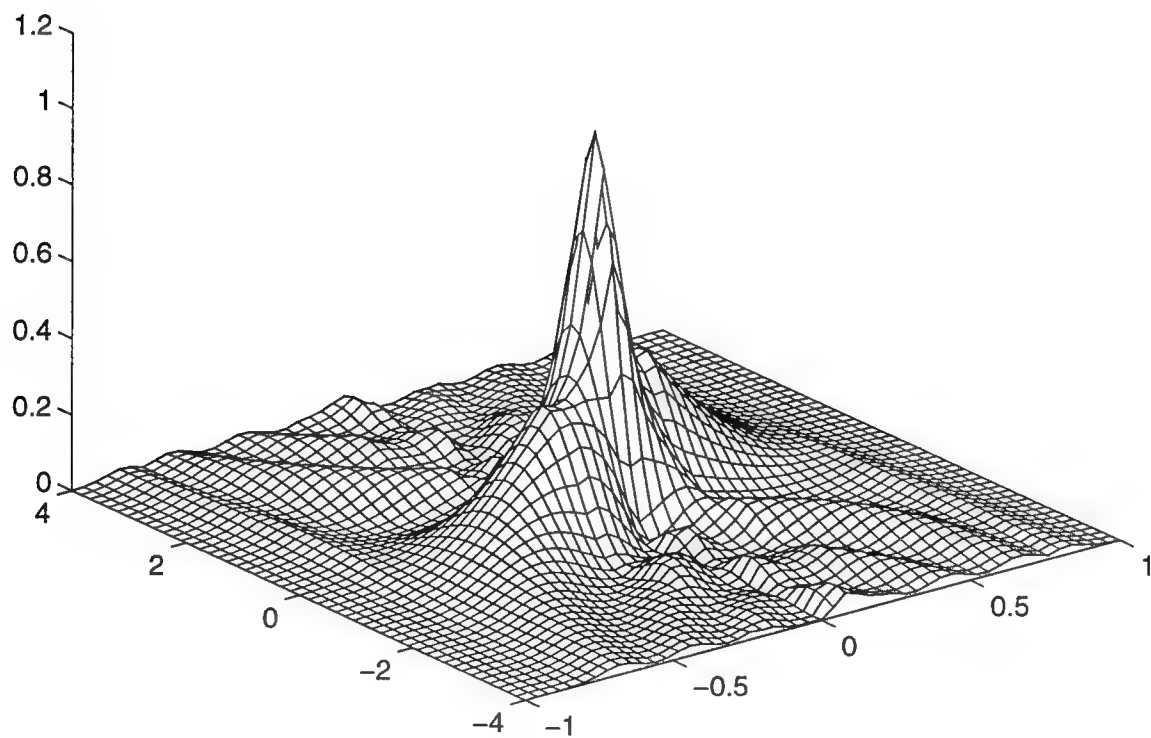
Figure 2.4: Ambiguity function for frequency sweep rate $k = 50\text{Hz/s}$: $\kappa = 30\%$.
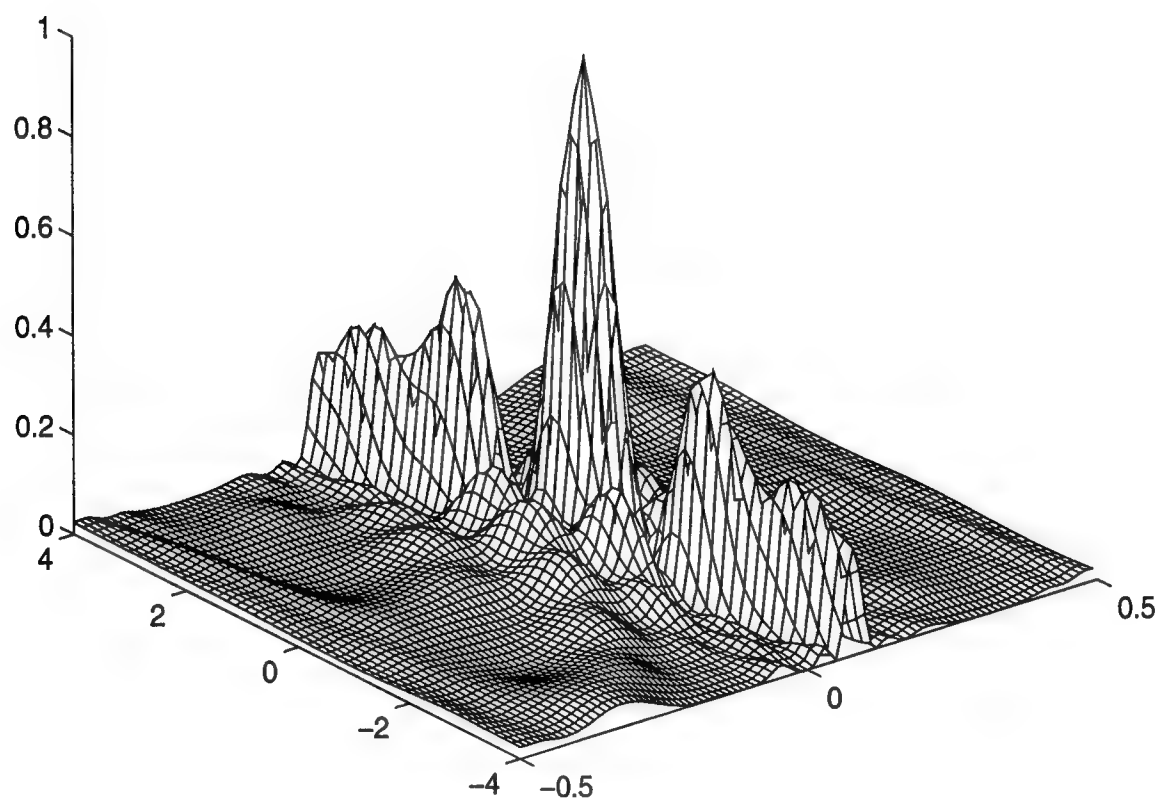
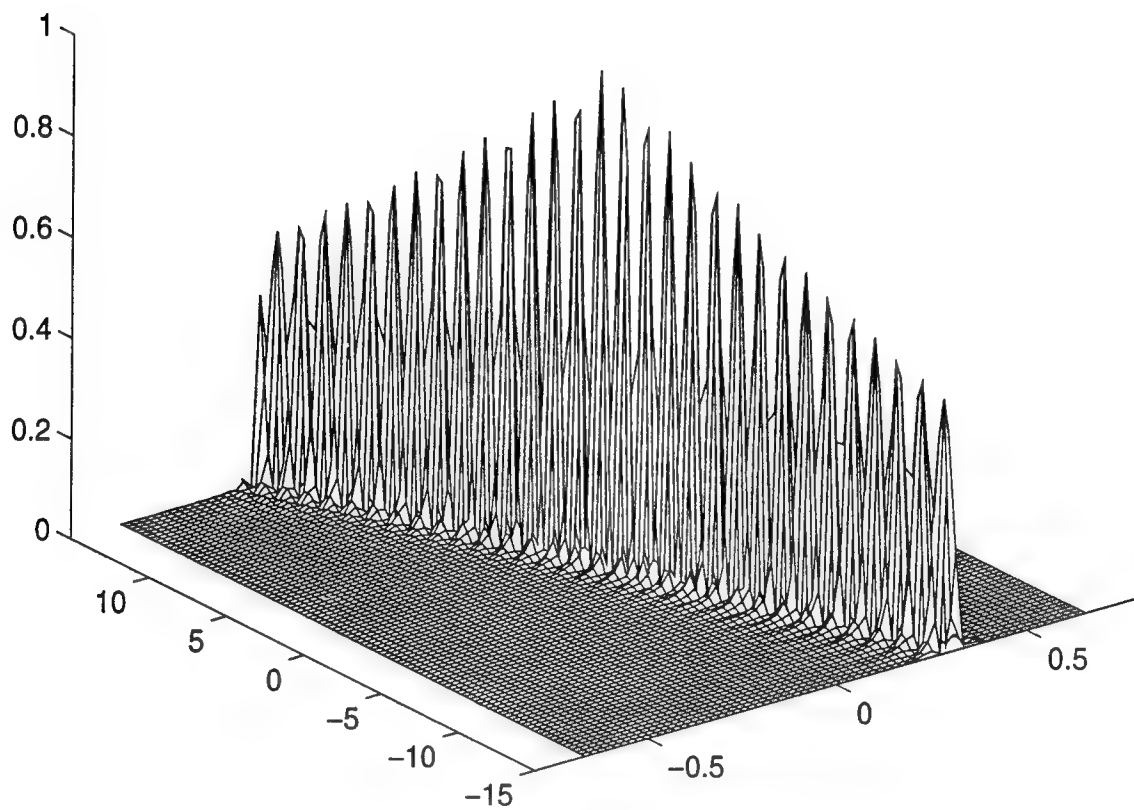Figure 2.5: Ambiguity function for frequency sweep rate $k = 50$ Hz/s: $\kappa = 60\%$.

Figure 2.6: Ambiguity function for frequency sweep rate $k = 50\text{Hz/s}$: $\kappa = 100\%$ (FM).
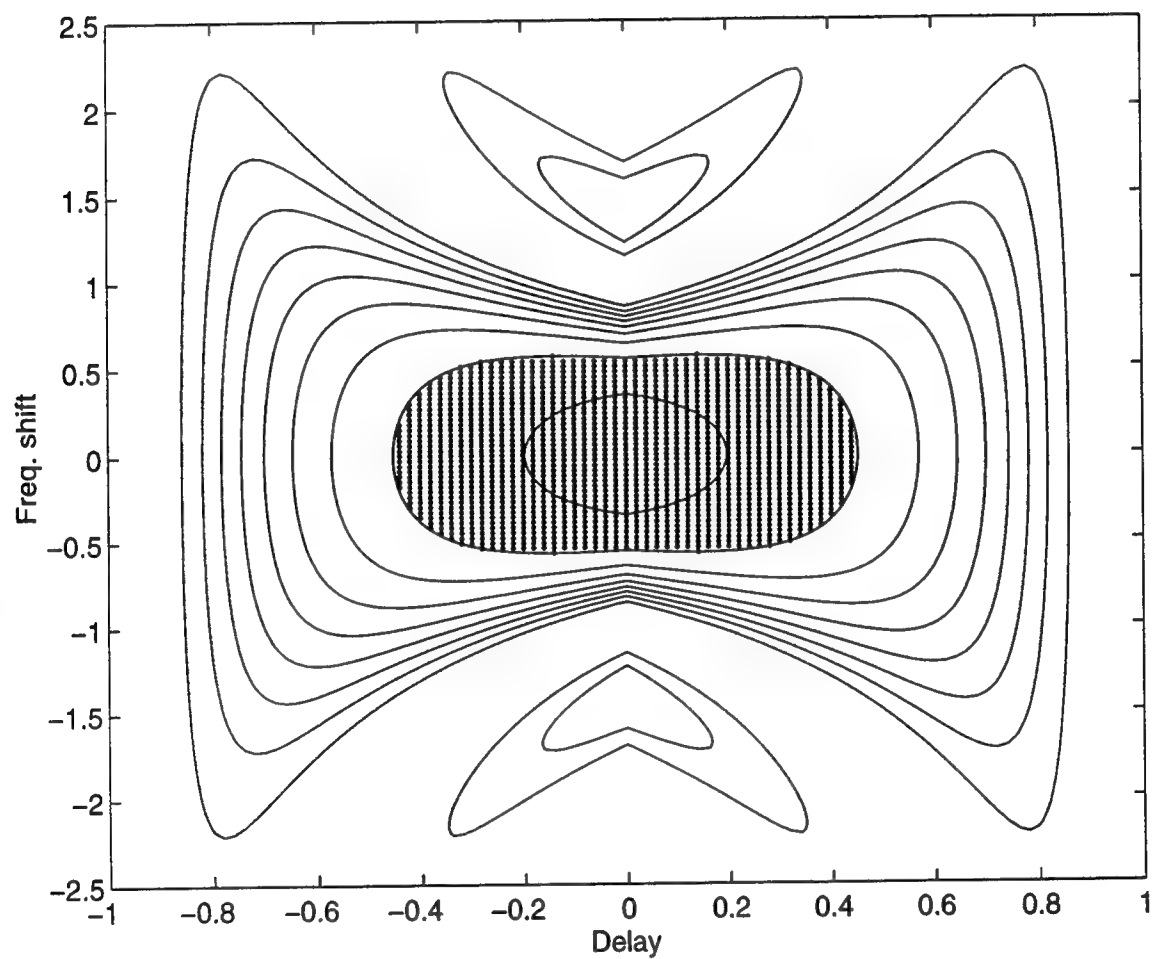
Figure 2.7: Probability of detection contours for $\alpha = .001$: $\kappa = 0\%$ (CW). In this case we have frequency sweep factor $k = 50\text{Hz/s}$.
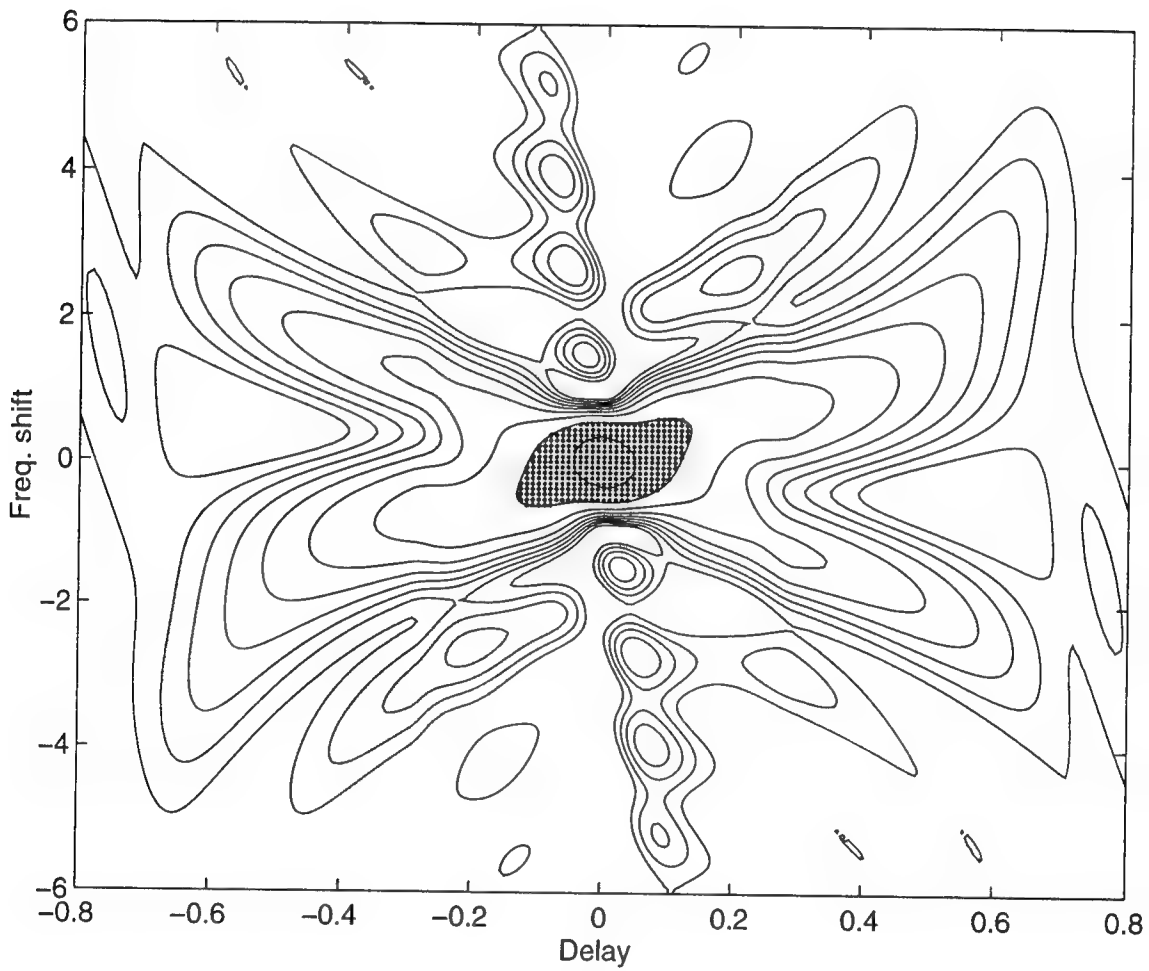
Figure 2.8: Probability of detection contours for $\alpha = .001$: $\kappa = 30\%$. In this case we have frequency sweep factor $k = 50\text{Hz/s}$.
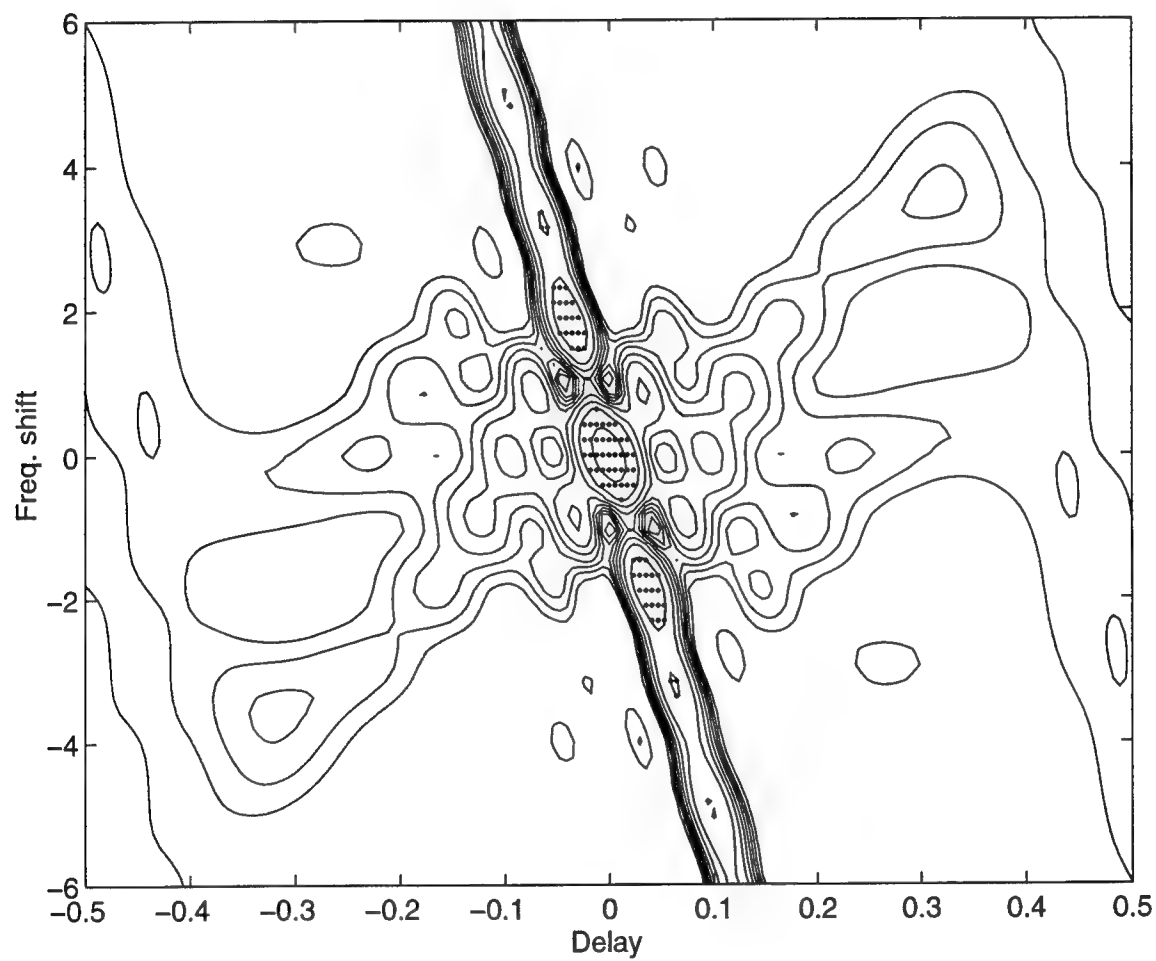
Figure 2.9: Probability of detection contours for $\alpha = .001$: $\kappa = 60\%$. In this case we have frequency sweep factor $k = 50\text{Hz/s}$.
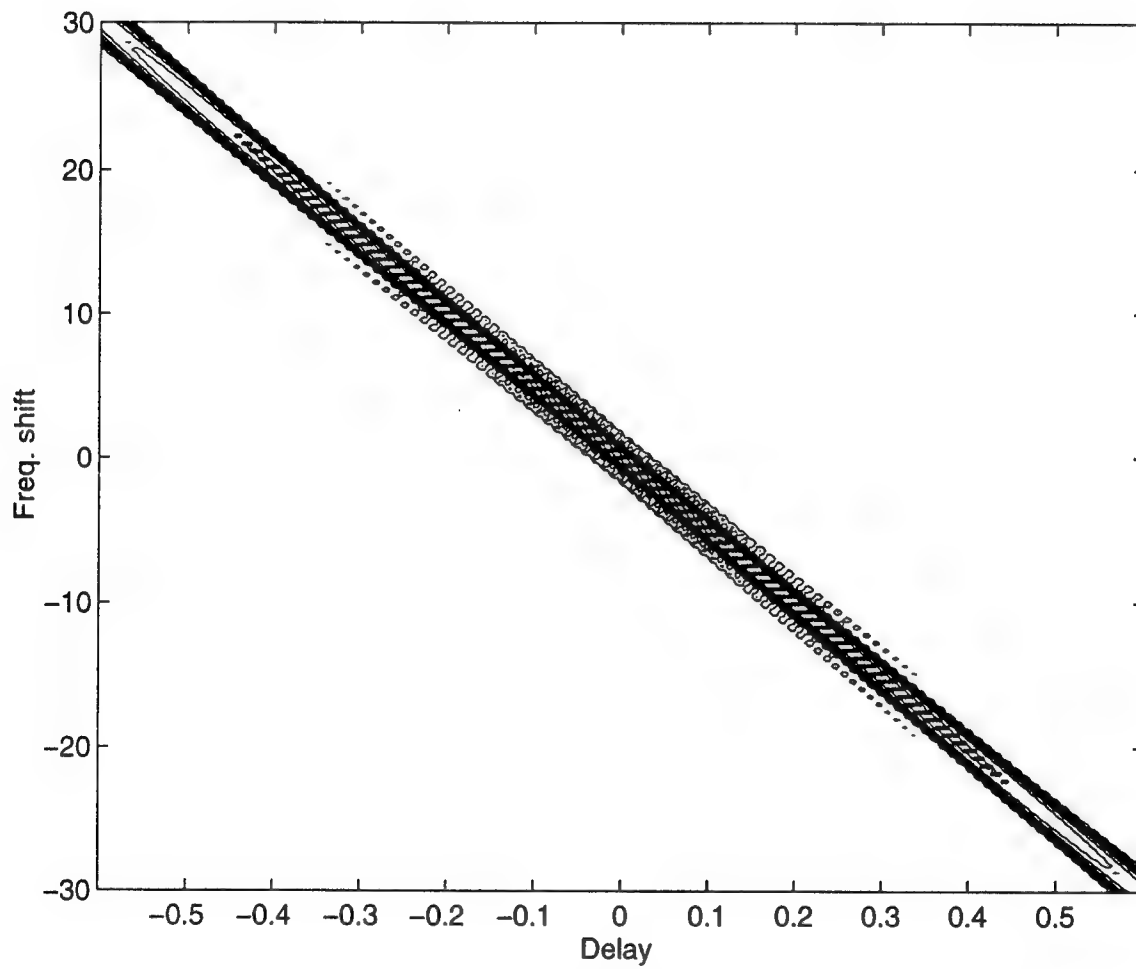
Figure 2.10: Probability of detection contours for $\alpha = .001$: $\kappa = 100\%$ (FM). In this case we have frequency sweep factor $k = 50\text{Hz/s}$.
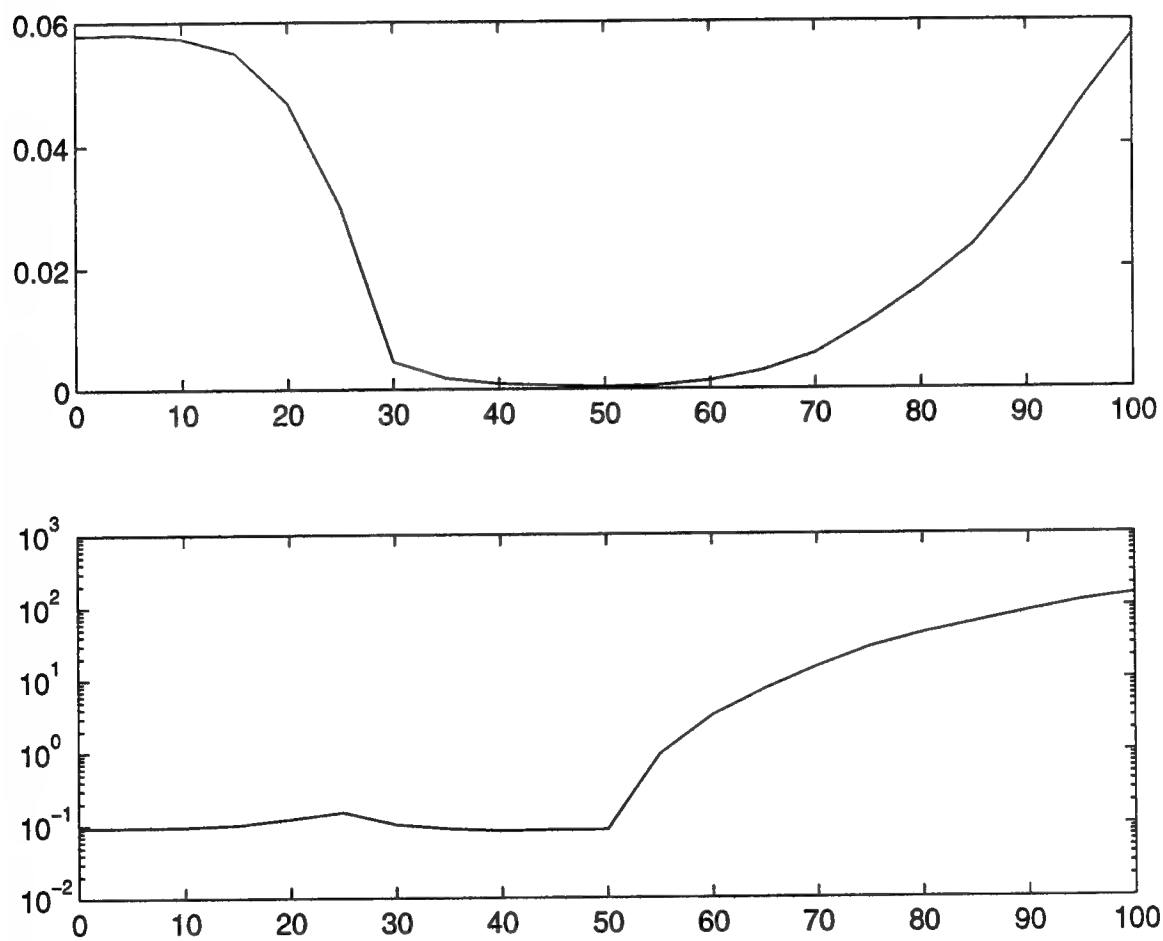
Figure 2.11: Measurement error variances as a function of κ, same parameters as figures 2.7, 2.8, 2.9, and 2.10. *Above:* Delay variance; *below:* Doppler-shift variance.
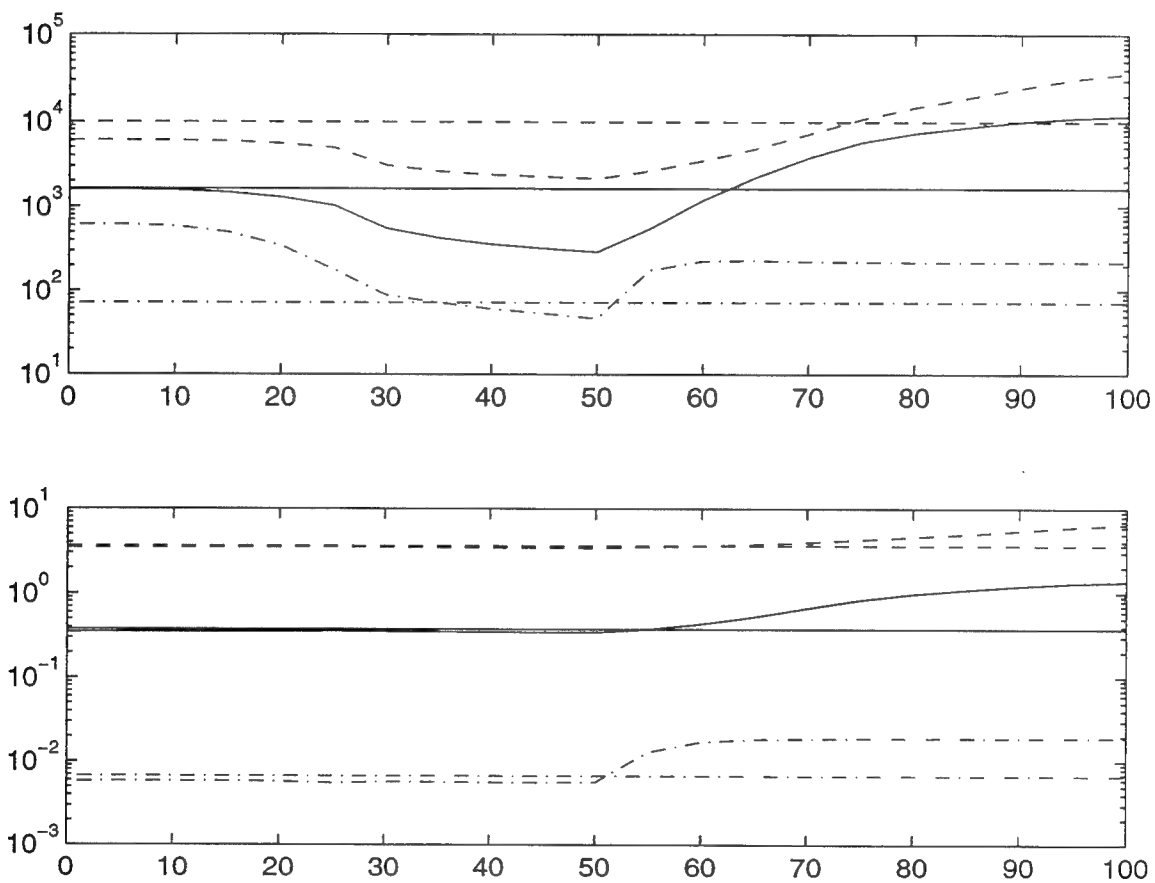
Figure 2.12: Steady-state estimation errors as a function of κ, same parameters as figures 2.7, 2.8, 2.9, and 2.10. *Above:* range variance; *below:* range-rate variance. Dashed lines: $\sigma_q^2 = 0.1$; solid lines: $\sigma_q^2 = 0.01$; Dash-dot lines: $\sigma_q^2 = 0.0001$. For each case the horizontal lines refer to the case of alternating CW/FM pulses.

Figure 2.13: Steady-state estimation errors as a function of κ, same parameters as figure 2.7, 2.8, 2.9, and 2.10, except that here the resolution cell is determined by the $P_d = 0.6$ contour. *Above*: range variance; *below*: range-rate variance. Dashed lines: $\sigma_q^2 = 0.1$; solid lines: $\sigma_q^2 = 0.01$; Dash-dot lines: $\sigma_q^2 = 0.0001$. For each case the horizontal lines refer to the case of alternating CW/FM pulses.
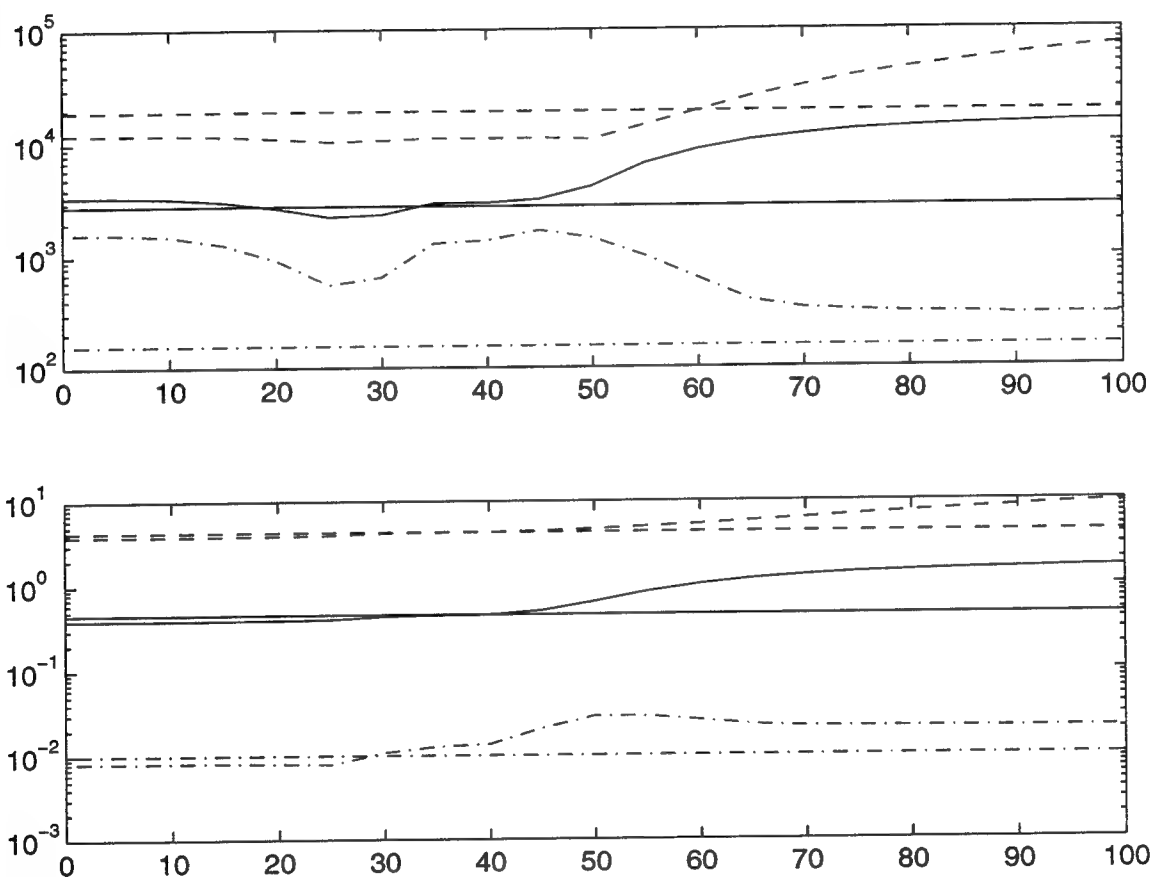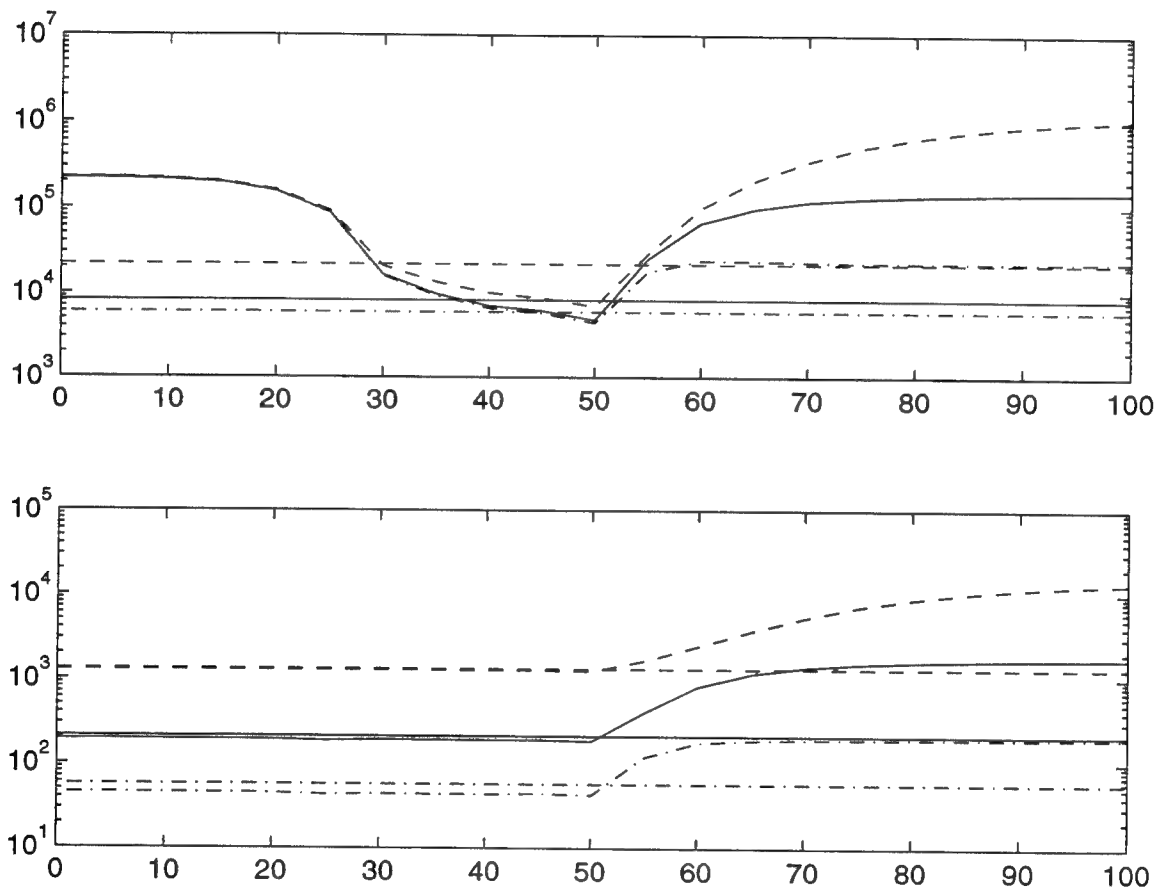
Figure 2.14: Steady-state estimation errors as a function of κ for the radar case described in section 2.4.2. *Above*: range variance; *below*: range-rate variance. Dashed lines: $\sigma_q^2 = 1000$; solid lines: $\sigma_q^2 = 100$; Dash-dot lines: $\sigma_q^2 = 10$. For each case the horizontal lines refer to the case of alternating CW/FM pulses.

Bibliography

[1] Yaakov Bar-Shalom and Thomas E. Fortmann, *Tracking and Data Association*, Academic Press, 1988.

[2] Tsitsiklis, J. (1993) Decentralized Detection.
In *Advances in Statistical Signal Processing*, Vol 2, pp.297-344, JAI Press.

[3] A. Elias-Fusté, A. Broquetas-Ibars, J. Antequera, and J. Yuste, "CFAR data fusion center with inhomogeneous receivers", *IEEE Transactions of Aerospace and Electronic Systems*, January 1992.

[4] Nadav Levanon, *Radar Principles*, John Wiley & Sons, 1988.

[5] Constantino Rago, Peter K. Willett, and Yaakov Bar-Shalom, "Resolution Cell Based Fusion – Making a Virtue of Necessity", *to appear in the Proceedings of the 1994 American Control Conference*, July 1994.

[6] Constantino Rago, Peter K. Willett, and Yaakov Bar-Shalom, "Tracking with fused noncoincident measurements", *Proceedings of the 1994 SPIE Conference 2235*, July 1994.

[7] J. Reeves and E. Reedy, *Principles of Modern Radar*, Van Nostrand Reinhold, 1987.

[8] X. Rong Li and Yaakov Bar-Shalom, "Stability Evaluation and Track Life of the PDAF for Tracking in Clutter", *IEEE Transactions on Automatic Control*, Vol. 36, No. 5, pp. 588-602, May 1991.

[9] Y. Bar-Shalom, X.R. Li, *Multitarget-Multisensor Tracking: Principles and Techniques*, YBS Publishing, 1995.

[10] Seymour, Stein, "Algorithms for Ambiguity Function Processing", *IEEE Transactions on Acoustic, Speech and Signal Processing*, ASSP-29, 3 (June 1981), 588-599.

[11] Y. Bar Shalom and X.R. Li, *Estimation and Tracking: Principles, Techniques and Software*, Artech House, Inc., 1993.

[12] C. Rago, P. Willett, and M. Alford, "Pre-Detection Fusion with Non-colocated Sensors: The Effect of the Resolution Cell Grid", *submitted to IEEE Transactions on Aerospace and Electronic Systems*, November 1994.

[13] C. Rago, P. Willett, and Y. Bar-Shalom, "Fusion Exploiting Overlap Amongst Resolution Cells", *Proceedings of the 1994 Dual-Use Technologies and Applications Conference*, Utica NY, July 1994.

[14] C. Rago, P. Willett, and M. Alford, "Pre-Detection Fusion with Non-colocated Sensors: The Effect of the Resolution Cell Grid", *University of Connecticut Technical Report TR-95-05*, (38 pages). April 1995.

[15] C. Rago, P. Willett, and Y. Bar-Shalom, "Combined CW-FM Pulse Fusion Schemes", *Proceedings of the 1995 SPIE Conference on Aerospace Sensing*, San Diego CA, July 1995.

[16] C. Rago, P. Willett, and Y. Bar-Shalom, "Direct Fusion of Constant- and Swept-Frequency Pulses", *in preparation for submittal to IEEE Transactions on Aerospace and Electronic Systems*.

Acknowledgement

MISSION

OF

ROME LABORATORY

Mission. The mission of Rome Laboratory is to advance the science and technologies of command, control, communications and intelligence and to transition them into systems to meet customer needs. To achieve this, Rome Lab:

    a. Conducts vigorous research, development and test programs in all applicable technologies;

    b. Transitions technology to current and future systems to improve operational capability, readiness, and supportability;

    c. Provides a full range of technical support to Air Force Materiel Command product centers and other Air Force organizations;

    d. Promotes transfer of technology to the private sector;

    e. Maintains leading edge technological expertise in the areas of surveillance, communications, command and control, intelligence, reliability science, electro-magnetic technology, photonics, signal processing, and computational science.

The thrust areas of technical competence include: Surveillance, Communications, Command and Control, Intelligence, Signal Processing, Computer Science and Technology, Electromagnetic Technology, Photonics and Reliability Sciences.